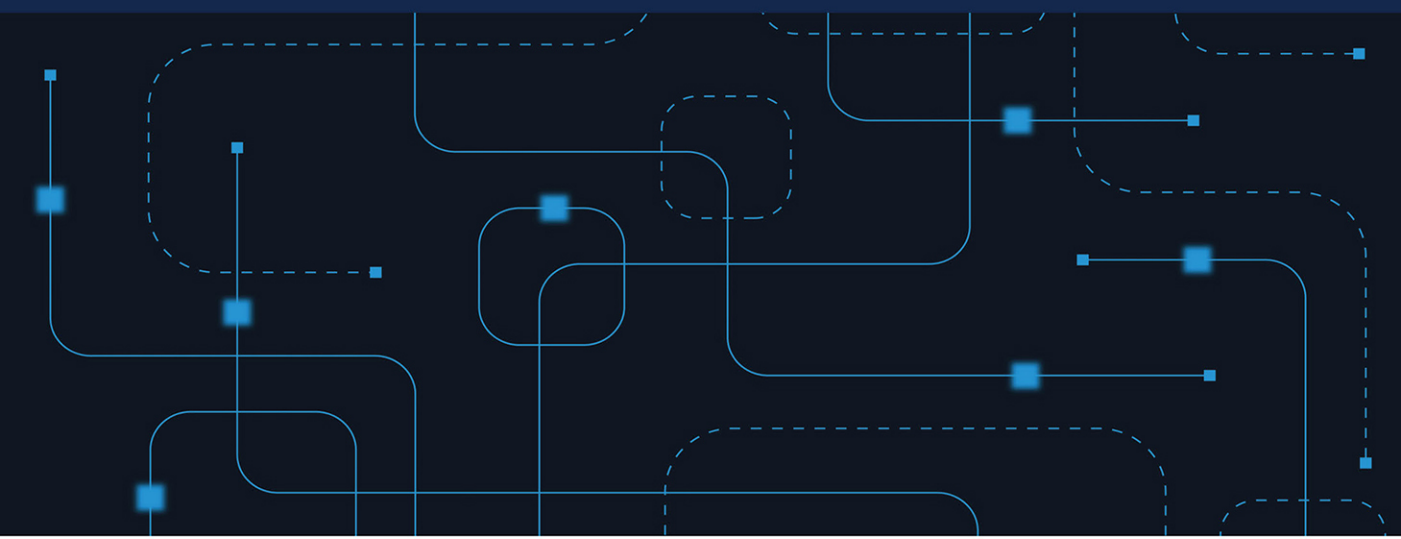
The Segmentation Blueprint

Designing, Deploying, and Scaling Network Segmentation for Modern Cyber Defense



RAGHUNATH KULKARNI, CISSP® No. 580752
KAARTHIK SIVAKUMAR
RENATO MORAIS, CCIE® No. 29902
PATRICK LLOYD, CCIE® No. 39750

ciscopress.com

FREE SAMPLE CHAPTER |

The Segmentation Blueprint

Designing, Deploying, and Scaling Network Segmentation for Modern Cyber Defense

Raghunath Kulkarni,
Kaarthik Sivakumar,
Renato Morais, CCIE No. 29902
Patrick Lloyd, CCIE No. 39750

Cisco Press

The Segmentation Blueprint: Designing, Deploying, and Scaling Network Segmentation for Modern Cyber Defense

Raghunath Kulkarni, Kaarthik Sivakumar, Renato Morais, and Patrick Lloyd

Copyright © 2026 Cisco Systems, Inc.

Published by:
Cisco Press
Hoboken, New Jersey

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

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

$PrintCode

Library of Congress Control Number: 2026935685

ISBN-13: 978-0-13-546236-2
ISBN-10: 0-13-546236-3

Warning and Disclaimer

This book is designed to provide information about segmentation strategies that the author team and technical team have implemented firsthand with various customers throughout their collective almost century-long tenure in the networking and security industries. This book covers the span from mindset, through planning, collecting data, strategizing, designing technology to align with strategy, through execution and implementation of segmentation strategies across domains. Every effort has been made to make this book as complete and as accurate as possible, but no warranty or fitness is implied.

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

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

Trademark Acknowledgments

All terms mentioned in this book that are known to be trademarks or service marks have been appropriately capitalized. Cisco Press or Cisco Systems, Inc., cannot attest to the accuracy of this information. Use of a term in this book should not be regarded as affecting the validity of any trademark or service mark.

Feedback Information

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

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

We greatly appreciate your assistance.

Head of IT & Professional Learning, Enterprise Learning and Skills: Julie Phifer

Alliances Manager, Cisco Press: Caroline Antonio

Executive Editor: James Manly

Managing Editor: Sandra Schroeder

Development Editor: Ellie C. Bru

Senior Project Editor: Tonya Simpson

Copy Editor: Kitty Wilson

Technical Editor: Luis Matos

Cover Designer: Chuti Prasertsith

Composition: codeMantra

Indexer: Brad Herriman

Proofreader: Donna E. Mulder



Americas Headquarters
Cisco Systems, Inc.
San Jose, CA

Asia Pacific Headquarters
Cisco Systems (USA) Pte. Ltd.
Singapore

Europe Headquarters
Cisco Systems International BV Amsterdam,
The Netherlands

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco Website at **www.cisco.com/go/offices.**

Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL: www.cisco.com/go/trademarks. Third party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1110R)

About the Authors

Raghunath Kulkarni is a principal engineer for the Cisco Security Technical Assistance Center (TAC) team. He focuses on cybersecurity operations and strategic collaboration, including threat intelligence, defensive measures, and ensuring that customers derive maximum value from their security investments. He has worked in technology and security at Cisco for over 16 years, ranging from deep technical support to cross-functional leadership roles with engineering and sales teams. His focus is guiding organizations through the adoption of robust defensive frameworks to minimize business risk and enhance operational resilience. Raghunath has extensive experience in advancing cybersecurity education across various sectors, including government, higher education, and enterprise security. Utilizing his CISSP and all five MITRE ATT&CK Defenders (MAD) certifications, Raghunath has architected comprehensive training programs and specialized engineering curricula for multiple universities. Raghunath's technology focuses include cyberthreat intelligence for proactive defense, testing and evaluation for security validation, and implementation of network security architectures to support digital transformation. Raghunath resides in Bengaluru, India, and when he is not consumed with technology, he is an avid reader and movie buff.

Kaarthik Sivakumar is a principal engineer in the Cisco Security Business Group, working in the Multi-Cloud team that comprises the Secure Workload and Isovalent products. He focuses primarily on microsegmentation platforms and policies. Kaarthik has been working in the industry for 26 years, ranging from routing protocols, forwarding tables, DWDM optical line system, and Cisco Secure Firewall. Kaarthik has spent many years in product security and trusted network infrastructures, such as using Trusted Platform Module for storing confidential material and validating the trust status of a device, topics on which he has published blogs on the Cisco Blogs platform. He has also published in IEEE journals on the security of microservices architectures and static analysis security testing. As an ACM India Eminent Speaker, he is invited to lectures in ACM chapters around the country on the topic of security and trust. Kaarthik resides in Chennai, India. He contributes some of his free time to open-source projects and tools to improve his personal life, some to his study of the Sanskrit language, and the rest to outdoor activities.

Renato Morais has been a cybersecurity solutions engineer at Cisco since 2017, and he has 20 years of experience in the networking and cybersecurity market. He works on projects with customers throughout the Americas, developing security architectures and solutions to protect users, data, and applications. A graduate in computer engineering, he holds various professional certifications, including CCIE, CISSP, and CCSP. Renato has also dedicated part of his career to training students in Cisco Networking Academy and preparing them for the CCNA certification. Furthermore, he has been recognized as a top contributor in the Cisco community for his contributions as a speaker at regional and global events such as Cisco Live and Cisco Connect, as well as for his work as a blog author.

Outside work, Renato enjoys home automation, music, books, movies, and sports, and he describes himself as a craft beer hunter. He lives in Campinas, Brazil, with his wife and two sons.

Patrick Lloyd is a senior solutions architect for the Cisco Security Professional Services team. He focuses on identity and access management, including segmentation, network access control, identity exchange, and identity integration for "smart" architectures in the continental United States and Canada region. Patrick has worked in technology delivery at Cisco for 16 years, including stints in the Technical Assistance Center (TAC) and working as a routing and switching design engineer, security design engineer, and security solutions architect. His focus is guiding customers through introducing structured approaches to increase visibility and identity exchange to minimize business risk and lateral attack vectors.

Previously, Patrick worked in higher education and defense industries in system administration and operational roles. Patrick has extensive experience in integrating identity into various industries, including healthcare, manufacturing, finance, and defense. Utilizing Cisco technologies and the methodologies covered in this book to build a layered security model, Patrick has architected segmentation architectures, including smart building architectures, for more than 100 customers. Patrick's technology focus is on TrustSec for segmentation, analyzing traffic flow with Cisco Secure Network Analytics/Stealthwatch for development of segmentation policies, implementing firewall and remote access architectures, and securing critical building systems through policy and segmentation while maintaining availability. Patrick resides in Durham, North Carolina, where he teaches self-defense and is an instrument-rated private pilot when not consumed with technology.

About the Technical Reviewer

Luis Matos is a principal architect on the Cisco CX Advanced Services team for the LATAM region, bringing more than 26 years of experience in the security field, with 23 years also spent as a university professor at various national institutions. He is a specialist in complex solutions for data centers, Internet/MPLS backbones, and security, with a strong focus on the telecommunications sector, large banks, and airports. Luis was the technical lead for the security team during the Rio 2016 Olympic Games and served as a security consultant for the 2015 Pan American Games in Toronto. Luis is a sought-after speaker who frequently presents technical subjects at Cisco events and IEEE events, in interviews with leading broadcast TV, at press companies across LATAM, and in lectures at universities. He holds multiple technical certifications in networking and security, including five CCIE certifications (Enterprise, Service Provider, Security, Voice, Data Center), DEVNET, CCNA, CCDA, CCNP, IPv6 Gold, MCP, MCP+INTERNET, and MCSE. In addition to his role at Cisco, Luis actively engages with various institutions and organizations to share his expertise. He has also conducted cybersecurity career presentations for public schools and for Cisco Networking Academy.

Dedications

To my parents: Words cannot fully express my gratitude for your endless encouragement, wisdom, and faith in me. Your values and guidance have been with me at every step, quietly shaping who I am and how I approach the world. Without your constant support and belief, this accomplishment would have remained a distant dream.

To my wife, Apurva Kittur: Thank you for being my rock. Your patience, understanding, and unwavering support throughout this journey have meant everything to me. In moments of doubt or difficulty, your steady presence and constancy have been my anchor.

To my daughter, Stuti Kulkarni: You are a daily reminder of the wonder of learning and the joy that curiosity brings. Even at just three years old, your laughter and boundless energy have inspired and motivated me more than you know.

This book is as much yours as it is mine. Thank you for being my greatest source of strength and for always being there for me.

—Raghunath Kulkarni

This book is dedicated to my parents, who taught me the focus and discipline to finish what I start. I also dedicate it to my wife and children, for their constant support. Finally, cheers to all the people who have the patience and strength to bring their ideas to life.

—Kaarthik Sivakumar

I dedicate this book to my grandfather, Luiz (in memoriam), who inspired my passion for reading and taught me to always pursue professional growth and excellence. To my mother, Maria, who invested in my education. To my wife, Chenia, who is by my side through every challenge. Above all, to my sons, Tomás and Bruno, who are my reasons for it all.

—Renato Morais

By the time this book comes to market, I'll have married the hardest-working and most driven person I've had the opportunity to know. For Courtney, my love and wife. The person who has had the patience while I work on another book while balancing work, flying, moving, and assisting in planning a destination wedding. Here's to new adventures in this crazy thing we call life.

Also, for Luna and Joey. While they'll never get to read this, their motivation and occasional needed distraction while deep in the throes of writer's block contributed to the success of my efforts culminated here. May we all have the joy and excitement they show for this adventure.

—Patrick Lloyd

Acknowledgments

Raghunath Kulkarni:

First and foremost, I would like to acknowledge Professor Shivanagowda G.M., under whose guidance I completed seven courses during my engineering studies in 2005. A particular statement from his lecture—"I am a proud owner of a copy of the reference textbook for this subject"—profoundly influenced me and inspired the aspiration to author a book that would be valued by the technical community. Even after nearly two decades, I recall the clarity and impact of his teaching methods, especially the analogies and examples that brought complex concepts to life. His commitment to education laid the foundation for this endeavor.

I am deeply grateful to my leadership team, and especially to Srinivasan Natarajan, for their steadfast support in fostering my passion for sharing knowledge. The open and supportive environment they cultivated was essential to undertaking and completing this project.

I also wish to thank my mentors—Subodh Gajare, Ankush Arora, and Ramses Smeyers—for their continued guidance and encouragement throughout my career and for supporting the vision of this book. My appreciation extends to my colleagues and friends, Sumit Bist, Avinash Naduvath, and Prashant Joshi, for their invaluable insights and willingness to engage in discussions that enriched the development of this work.

To the engineers in TAC, who spend countless hours supporting our customers: Your strong technical acumen has kept me grounded when some of my ideas were as imaginative as warp drive. Sadhana K.S., Ravi Shankar K.B., and Nikhil Alampalli Ramu, you have been exceptional.

Finally, I would like to acknowledge my co-authors, whose professionalism and collaborative spirit have been central to the success of this project. Our collective dedication, despite geographic distances, enabled us to refine and unify diverse perspectives. It has been a privilege to work alongside such talented individuals.

Kaarthik Sivakumar:

I want to thank the many managers throughout my career who gave me the freedom to explore beyond my day-to-day work—to experiment, learn something new every day, teach, and bring my ideas to life, even when there were always fires at work to put out. My thanks go especially to my current manager, Jana Radhakrishnan, who always trusted me to do the right thing and encouraged me to write this book.

I am grateful to my team, who taught me a great deal while making it look like I was the one teaching them. I also want to thank Alex Levin, Brijeshkumar Shah, Jorge Quintero, Abhijeet Nakhate, Jerome Blomart, Furong Gisiger, Aamir Ahmed, Sam Kilaru, and Srinivasa Rao Pisupati, for showing me how our customers use our products. Without that insight, this book would never have come to life.

Last but most definitely not least, I want to thank my co-authors for being a part of this journey, especially on my first book. I hope we can do this together many more times.

Renato Morais:

In 2005, I learned about this company called Cisco, and it was love at first sight. That moment marked the beginning of a professional journey where dreams have turned into goals and goals into achievements. Ideas that once seemed distant, or even impossible, became reality. When I achieved my CCNA, I thought, "Imagine one day becoming a CCIE." I used to attend Cisco webinars; today, I am a regular speaker at them. There was a time when I dreamed of going to Cisco Live, and now I can say that I am a speaker. Each new official book I acquired was a source of inspiration, and now I can see my own name on a cover.

Along this journey, I have had the privilege of working with and learning from brilliant and inspiring people.

I am grateful to the mentors who have guided me throughout my career: Cláudio Silva, who set me on the path to becoming a Cisco specialist and Networking Academy instructor; Alberto Oliveira, who inspired my pursuit of professional excellence; and Luis Matos and Marcello Ishida, colleagues who have helped guide my way within Cisco.

I am grateful to my friend and brother-in-arms on this journey, Bruno Fagioli, who has been by my side since our CCIE study days and shares my passion for challenges and the pursuit of knowledge.

To my co-authors, thank you for accepting this challenge and making the realization of this dream possible.

And I am grateful to my wife and sons, who are always by my side and motivate me to keep moving forward.

Patrick Lloyd:

The writing team on my previous book, *Zero Trust Architecture*, spent countless hours together, both physically and virtually, as we navigated our extensive careers. In 2017, I recruited David Houck into Cisco Security Services, seeing the potential, drive, and empathetic nature that he brought to customers while working for Cisco TAC. In his dedication he wrote, "For all who teach, inspire, and mentor us." In intersection with his phrasing, my acknowledgment in that work was meant to include people like David who shared the passion for knowledge sharing, mentorship, and a scrupulous attitude to benefit the industry and our customers: "In a technology career spanning multiple decades, one does not reach the apexes in technical knowledge without exposure and guidance from some of the best managers, mentors, and sounding boards in the industry." The loss of his knowledge, empathy, charity, and drive can never be replaced but can only serve as a guiding light for us to aspire to.

Huge thanks to my management team, Chris Mula in particular, who has never pushed back on my aspirations to continue to share knowledge in a public forum through works such as this, industry presentations, and webinars.

Thank you to my engineers—the ones who keep me honest, push back on some of my vague and lofty ideas, and spend countless hours supporting our customers in their goals: Ryan Wehe, Chris Jones, and Logan Helfrich.

Thanks also to my mentors and friends who have continued to support my career and aspirations to develop the knowledge demonstrated in this book and who encourage me to expand on it day after day: Jon Taylor, Dave Lucas, and Cindy Green-Ortiz.

Finally, thanks to this team of authors. Our relationship has always been a virtual one and started with the ask of whether I had feedback on a book outline they aspired to write on designing segmentation architectures. Feedback turned into an offer to help in writing this work, and I'm thankful to them for allowing me to continue to share my experience with the technology community.

Contents at a Glance

Introduction xxiii

Chapter 1 The Segmentation Mindset 1

Chapter 2 Alignment with Business Outcomes 25

Chapter 3 Developing a Segmentation Strategy 71

Chapter 4 Macrosegmentation 117

Chapter 5 Microsegmentation 147

Chapter 6 Building the Segmentation Fabric 189

Chapter 7 Implementing Segmentation with Cisco Technology 231

Chapter 8 Segmenting Applications in the Data Center and Cloud 275

Chapter 9 Validating Policies, Monitoring Enforcement, and Responding to Deviations 317

Chapter 10 Segmentation Maturity Model and Scorecard 347

Chapter 11 Reference Architecture 395

Chapter 12 The Future of Segmentation 437

Appendix: Leadership Perspective Scorecard Approach 453

Index 461

Reader Services

Register your copy at www.ciscopress.com/title/9780135462362 for convenient access to downloads, updates, and corrections as they become available. To start the registration process, go to www.ciscopress.com/register and log in or create an account.* Enter the product ISBN 9780135462362 and click Submit. When the process is complete, you will find any available bonus content under Registered Products.

*Be sure to check the box that you would like to hear from us to receive exclusive discounts on future editions of this product.

Contents

Introduction xxiii

Chapter 1 The Segmentation Mindset 1

Pillars of Zero Trust 2
- Policy and Governance 2
- Identity 5
- Vulnerability Management 8
- Enforcement 9
- Analytics 10

Beyond the Five Pillars 11

Addressing the Problem Head On: Why Do Segmentation Strategies Fail? 11
- The Importance of the Team 13
- Information Security 14
- Network Engineering/Desktop Engineering 15
- Network Security 15
- Operations 16
- Cross-Team Collaboration 17
- Business Entities 18
- Executive-Driven Vision 19
- Other Considerations Beyond Standard Teams 19

Aligning Strategy and Tactics 21

Maturing in Segmentation Strategy 22

A Look Ahead 23

Summary 24

Reference 24

Chapter 2 Alignment with Business Outcomes 25

Identifying the Need for Segmentation in Modern Hybrid Networks 25

Patterns Driving Modern Hybrid Networks 27
- Bring Your Own Device (BYOD) 27
- Internet of Things (IoT) 28
- Artificial Intelligence (AI) 29

Security Risks: A Hacker's Playground 30

Compliance Mandates: The Regulatory Maze 33
- Microsegmentation for Regulatory Compliance 34
- Payment Card Industry Data Security Standard (PCI DSS) 35

Health Insurance Portability and Accountability Act (HIPAA) 37
General Data Protection Regulation (GDPR) 39
Performance, Scalability, and Adaptability 42
Operational Complexity: The Management Nightmare 43
Cloud and Multicloud Adoption: The Consistency Crisis 45
Developing a Segmentation Strategy 45
Reducing Risk via Organizational Perspectives 46
The Leadership Perspective 46
Strategic Oversight 47
Effort Versus Value 48
Resource Allocation 50
Policy Development 51
Risk Management 52
Stakeholder Engagement 53
Training and Awareness 54
Performance Metrics 54
Continuous Improvement 55
The Architect's Perspective 56
Design and Planning 56
Integration and Alignment 58
Innovation and Adaptation 58
Collaboration and Communication 59
The Asset Management Perspective 60
The Network and Security Administrator's Perspectives 61
The Application Development Perspective 63
Perspectives on Infrastructure, Network, Applications, and Automation 64
Roles and Responsibilities of the Application Team for Effective Segmentation 65
Bringing It All Together: An Actionable Segmentation Framework 66
Summary 68
References 69

Chapter 3 Developing a Segmentation Strategy 71

Cisco SAFE 72
How to Use Cisco SAFE 73
Capability Phase 75

Architecture Phase 77
Design Phase 82
The Foundations of Segmentation 85
Physical Segmentation 86
Logical Segmentation 87
Virtual Local Area Networks (VLANs) 87
Private VLANs (PVLANs) 90
Wireless SSID 92
Access Control List (ACL) 92
SD-Segmentation and Cisco TrustSec 94
Security Zones 97
Network Virtualization 99
Extending Network Segments with VXLAN 101
Session Layer Segmentation with QUIC 102
Segmentation as a Service in Public Cloud 105
Cloud-Native Segmentation 106
Kernel-Level Segmentation with eBPF 112
Segmentation Strategy and the Shared Responsibility Model 113
Summary 115

Chapter 4 Macrosegmentation 117
Gaining Visibility While Architecting Segmentation 118
Network Virtualization 120
An Overlay for an Overlay: VLANs 122
Achieving Macrosegmentation with Firewalls 123
Understanding Traditional Access Control Lists 124
Understanding Interface-Based Firewalls 125
Understanding Zone-Based Firewalls 126
Hybrid Mesh Firewall 127
Linking Concepts Together 130
Practical Macrosegmentation Policy Development 130
Segmentation Involves the Infrastructure but Is Really About the Endpoint 134
Mapping with Cisco SAFE Architecture 136
Campus: Securing Guest Wi-Fi Access in a Campus Network 136
Solution with Macrosegmentation Features 137
Conclusion 137

Branch: Protecting Payment Processing in a Branch Network 137
A Solution with Macrosegmentation Features 138
Conclusion 139
Data Center: Securing Remote Access in a Data Center 139
Solution with Macrosegmentation Features 140
Conclusion 140
Internet PIN: Securing SaaS Application Access in the Internet PIN 140
Solution with Macrosegmentation Features 141
Conclusion 141
Cloud PIN: Securing Workload Communication in a Cloud PIN 142
Solution with Macrosegmentation Features 142
Conclusion 143
Edge PIN: Securing IoT Endpoints in an Edge PIN 143
Solution with Macrosegmentation Features 143
Conclusion 145
Summary 145
References 145

Chapter 5 Microsegmentation 147
Benefits of Microsegmentation 148
Implementing Microsegmentation 148
Challenges in Implementing Microsegmentation 150
Microsegmentation in the Campus 152
Application Segmentation 153
Cloud-Native Segmentation Controls 154
Organizing Workloads 155
Segmentation 156
Applying Policies 157
Network Service Mesh 160
Automated Zero Trust Microsegmentation 160
Grouping Workloads 162
Organizing Workloads 164
Automating Scope Discovery 166
Critical Common Services 169
Providing Access to Scopes 170
Workload Information 170

Policies 171
Policy Creation 172
Policy Discovery 174
Measuring Segmentation 175
Achieving Microsegmentation with a Next-Generation Firewall 175
Seeing Through the Fog: Application Awareness 176
Knowing Who's at the Gate: User Identity Policies 177
Labeling the Landscape: Security Group Tags 178
Reading the Room: Context-Aware Policies 179
Cisco Secure Firewall: The Skilled Artisan 179
Applying More Granular Enforcement Mechanisms Closer to the Endpoint 182
Integration with Other Cisco Technologies for Enhanced Segmentation 186
Summary 188
References 188

Chapter 6 Building the Segmentation Fabric 189
Cisco SD-Access Components 190
Cisco Catalyst Center 190
Cisco Identity Services Engine 192
Operational Planes 195
LISP: The Overlay Control Plane 195
VXLAN: The Data Plane 195
Cisco TrustSec: The Policy Plane 196
Cisco Catalyst Center: The Management Plane 196
Architecture Components 197
Fabric 197
Underlay Network 197
Overlay Network 198
Shared Services 200
Fabric Roles 201
Control Plane Node 201
Edge Nodes 203
Intermediate Nodes 204
Border Nodes 204
Extended Nodes 208
Fabric WLCs 209
Fabric-Mode Access Points 210
SD-Access Embedded Wireless 210

Fabric in a Box 210
Transit Networks 211
Transit Control Plane Nodes 211
Fabric Site 212
SD-Access Design Strategy 212
Small Sites 213
Medium Sites 214
Large Sites 215
A Reference Model to Deploy SD-Access for Distributed Campus 217
Interconnecting Multiple Domains to Enable End-to-End Segmentation over the Internet 219
End-to-End Segmentation with Cisco SD-WAN 220
Cisco SD-WAN Design Considerations 223
Cisco Secure Access: Cloud-Based Security Controls with Context-Aware Policy Enforcement 224
Implementing Secure Private Access with Cisco Secure Access 226
Delivering Secure Internet Access with Cisco Secure Access 228
Achieving Seamless Segmentation Using Integrated Cisco Solutions 229
Summary 230
References 230

Chapter 7 Implementing Segmentation with Cisco Technology 231
Configuring Cisco Secure Firewalls for Segmentation 231
Application Visibility and Control (AVC) 231
User Identity–Based Policies 233
Integrating with Identity Services Engine 235
Expanding Integrations to Include Cisco Secure Workload 237
Using Context-Aware Policies 237
Configuring Cisco Identity Services and Switches for Segmentation 240
Configuring a Switch to Communicate to ISE 244
Configuring a Dynamic VLAN 245
Configuring a Downloadable ACL 246
Configuring TrustSec Tags 247
Assigning Policies to the TrustSec Matrix 251
Configuring Policies 255
Building the Network Inventory with Endpoint Profiling 256
Cisco AI Endpoint Analytics: Advanced Endpoint Classification 259
Achieving Consistency with ISE Data 263

Simplifying Multidomain Segmentation with Cisco ISE Common Policy 263
Introduction to Cisco ISE pxGrid 264
SXP: Enabling SGT Propagation Across the Network 265
Cisco ACI Integration: Extending Control into the Data Center 266
Public Cloud Integrations: Mapping Workload Context for Unified Policy Enforcement 268
CMDB Integration: Mapping the Unknowns 269
Cisco Cyber Vision Integration: Enhanced Visibility for Industrial Networks 270
Third-Party Integrations: Extending Visibility and Control with Specialized Solutions 271
Ordr: Endpoint Visibility and Behavioral Insights 272
Medigate: Optimizing Security in Healthcare Environments 272
Armis: Comprehensive Asset Security Across IoT and OT 273
Summary 273
References 274

Chapter 8 Segmenting Applications in the Data Center and Cloud 275
Zero Trust Microsegmentation with Cisco Secure Workload 275
Cisco Secure Workload First-Time User Experience 276
Visibility into Network Traffic Flows 280
Workload Profiles 284
Policies 284
Automated Policy Discovery Using AI 287
Measuring Segmentation Scores 289
Segmenting the Data Center with Cisco ACI 290
Components of an ACI Fabric 291
Physical Components 291
Logical Constructs (Policy Model) 292
Network-Centric Versus Application-Centric Topology 293
Network-Centric Topology 294
Application-Centric Topology 294
ACI Fabric Overlay 295
External Connectivity with L3Out 296
Connectivity with an SD-Access Fabric 297
ACI Fabric Communication Through a Catalyst SD-WAN Mesh 298
Virtual Machine Manager (VMM) Integration 299
Container Network Interface (CNI) Integration 303

Service Graphs for Advanced Security Inspections 305
Network Segmentation in ACI 306
Cisco Secure Workload and ACI 309
The Life of Packets in the End-to-End Segmented Network 312
Summary 315
Reference 316

Chapter 9 Validating Policies, Monitoring Enforcement, and Responding to Deviations 317

Cisco Secure Network Analytics: From Visibility Gaps to Actionable Intelligence 318
Components and Architecture 320
Flow Ingestion 322
Host Groups 327
Observations, Alerts, Security Events, and Alarms 330
Custom Security Events 334
Group Policies Monitoring 335
Response Management 336
Technical Adoption Roadmap 339
Cisco Secure Workload: Policy Validation with Live Policy Analysis 340
Summary 344
References 345

Chapter 10 Segmentation Maturity Model and Scorecard 347

What Is the Segmentation Maturity Model? 348
Using the Capability Maturity Model (CMM) for Segmentation 348
Level 1: Ad Hoc/Initial 349
Level 2: Repeatable/Managed 349
Level 3: Defined 350
Level 4: Quantitatively Managed 350
Level 5: Optimized 350
Segmentation Methods and the OSI Model 351
Threat Assessments 352
Factors That Trigger a Maturity Level Reassessment 353
Role-Based Perspectives 353
The Management and Senior Leadership Perspective 354
Strategic Oversight 354
Resource Allocation 354
Policy Development 355

Risk Management 355
Stakeholder Engagement 356
Training and Awareness 356
Performance Metrics 357
Continuous Improvement 357
Evaluating Leadership Effectiveness: A Scorecard Approach 358
Mapping the Scorecard and Evaluation Criteria with the SSMM 360
The Asset Management Perspective 362
Asset Inventory Management 363
Lifecycle Management 363
End-User Device Management 363
Infrastructure Device Management 364
Evaluating Asset Management Effectiveness 365
Mapping the Scorecard and Evaluation Criteria with the SSMM 367
The Architect's Perspective 367
Design and Planning 368
Integration and Alignment 370
Innovation and Adaptation 372
Collaboration and Communication 374
Evaluating Architecture Effectiveness: A Scorecard Approach 375
Mapping the Scorecard and Evaluation Criteria with the SSMM 377
The Network and Security Administrator's Perspective 378
Network Administrator Responsibilities 379
Design and Implementation 379
Configuration and Maintenance 380
Monitoring and Optimization 380
Integration Across Hybrid Environments 380
Security Administrator Responsibilities 381
Security Policy Enforcement 381
Threat Detection and Response 381
Identity and Access Management 382
Incident Management 382
Evaluating Network and Security Administrative Effectiveness: A Scorecard Approach 383
Alternative Approach 383
Mapping the Scorecard and Evaluation Criteria with the SSMM 390
Summary 392

Chapter 11 Reference Architecture 395

The Campus and Branch Domains 399

Use Case Example: Smart Hospital, Inc. 411

Progressing Beyond an Endpoint Focus 415

A Ubiquitous Experience for All 417

Beyond Campus and Branch 419

Identifying Segmentation Needs with Cisco SAFE 424

Building a Secure Architecture with Cisco SAFE 429

Previous Reference Architectures 430

Asset Identity Management, Asset Monitoring and Discovery, Configuration Management Database, and a Sound Provisioning or Onboarding Process 431

Authentication, Authorization, and Accounting (AAA), Certificate Authorities (CAs), and IP Address Management (IPAM) Systems 432

Traffic Visibility, Behavioral Analytics, Firewalls, Proxies, and DNS Security 433

Analytics, Logging, and Lessons Learned 433

Summary 434

References 435

Chapter 12 The Future of Segmentation 437

Core Elements of Segmentation 437

The Future Role of AI in Segmentation 438

Zero Trust: The Guiding Principle for Granular Segmentation 440

Identity as the New Perimeter 440

Contextual Awareness for Adaptive Access 440

Continuous Verification and Least Privilege 442

Microperimeters and Nanosegmentation 443

Data Privacy as a Segmentation Imperative 444

A Vision for the Future of Segmentation: Distributed Macrosegmentation and Microsegmentation with a Hybrid Mesh Firewall 445

The Impact on Nanosegmentation 450

Automation: Streamlining Segmentation Processes 450

The Impact of Cloud-Native and On-Premises Technologies 451

A Call to Action: Integration of Technologies 451

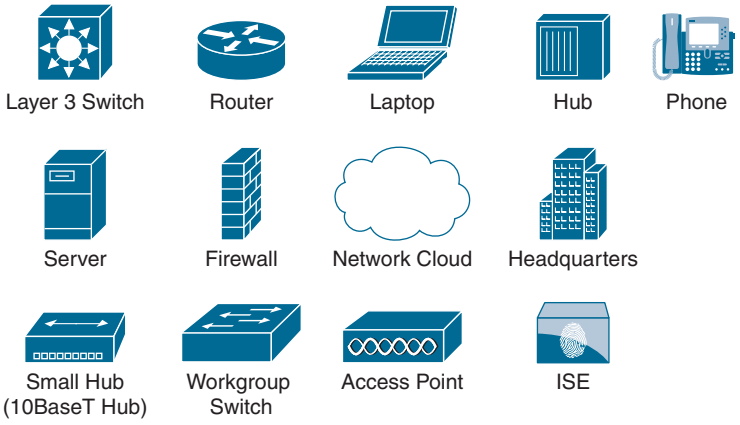
Summary 451

Reference 451

Appendix: Leadership Perspective Scorecard Approach 453

Index 461

Icons Used in This Book



Introduction

While the concept of segmentation is not new to network design, the increasing commonalities of lateral attack vectors, aging endpoints with limited updates made to software and firmware, and increased commonality of systems interacting with each other in undocumented ways drive a need for increased focus on macrosegmentation and microsegmentation. This book is meant to serve as a guiding light for best practices in implementing segmentation across domains, providing a starting point in terms of mindset, required stakeholders, and domains to be focused on to increase the security of enterprise networks. This book serves as a work on its own and also refers to concepts in various other works, including *Zero Trust Architecture*, which explores a broader approach to securing the network.

Who Should Read This Book?

This book is designed for network architects, security engineers, and IT strategists who recognize that in today's borderless digital landscape, the traditional perimeter is no longer sufficient. As lateral movement becomes the primary objective for modern attackers, segmentation has emerged as the foundational necessity for any Zero Trust strategy.

How This Book Is Organized

- **Chapter 1, "The Segmentation Mindset":** Traditional brownfield networks are often too complex for simple security fixes, and technical implementations frequently fail because they ignore the human elements of an organization. This chapter establishes the five-pillar Zero Trust model (policy, identity, vulnerability management, enforcement, and analytics) and addresses the critical organizational hurdles of politics, money, and culture that must be navigated to ensure project success.
- **Chapter 2, "Alignment with Business Outcomes":** For the modern enterprise, security is a board room priority. Cyberthreats now directly impact brand reputation, customer trust, and the bottom line, and IT is increasingly being pressured to move from being a cost center to being a strategic enabler. This chapter addresses the business imperatives for segmentation, focusing on risk management and regulatory adherence. It explores how to balance the security versus connectivity spectrum and how automation drives operational efficiency and cost reduction in the face of disruptive trends like BYOD, IoT, and AI.
- **Chapter 3, "Developing a Segmentation Strategy":** Organizations often struggle to translate high-level security goals into specific technical requirements without a repeatable methodology. This chapter introduces the Cisco SAFE architecture as a modular reference model to map business flows to specific security capabilities. Beyond the OSI model, it addresses architectural blueprints for different places in the network (PINs)—such as branches, campuses, and clouds—and defines operational secure domains, including compliance, security intelligence, and management.

- **Chapter 4, "Macrosegmentation":** Broad network isolation is the first step in containing a breach, but inconsistent naming and routing across sites often lead to "flat" networks where all segments can communicate without restriction. This chapter addresses the tactical implementation of virtual routing and forwarding (VRF) and zone-based firewalls.
- **Chapter 5, "Microsegmentation":** The true challenge of any segmentation architecture is implementing microsegmentation—due to the sheer amounts of data required to do so and the need to aggregate, organize, and inventory workloads and endpoints. This chapter explores these challenges and approaches to organize and understand workloads and endpoint interactions and to design microsegmentation policies to be applied to each. It explores requirements to create a successful microsegmentation strategy across multiple network focus areas—campus, data center, remote access, and cloud.
- **Chapter 6, "Building the Segmentation Fabric":** In this chapter, we roll up our sleeves and explain practical implementation steps for building a robust fabric capable of carrying layered segmentation methodologies. We explore various methods of designing the fabric and placing devices that hold various roles within the fabric. We explain domains that must interact, such as multiple campuses, and how the WAN lends itself to a segmentation strategy. This chapter sets the stage for future chapters on specific enforcement mechanisms for the campus and data center environments that can be applied based on use case.
- **Chapter 7, "Implementing Segmentation with Cisco Technology":** This chapter explores practical concepts, such as various segmentation methods that can be applied throughout the architecture and their strengths within the security onion. It presents configuration steps for implementing many methods and walks through implementation steps that can be applied for each of the technologies.
- **Chapter 8, "Segmenting Applications in the Data Center and Cloud":** This chapter discusses the application of segmentation policies based on profiling workload behaviors using Cisco Secure Workload (CSW) and Cisco Application Centric Infrastructure (ACI) for the data center and hybrid cloud. It breaks down components of a segmentation fabric into their respective parts and discusses how they integrate with an SD-WAN mesh.
- **Chapter 9, "Validating Policies, Monitoring Enforcement, and Responding to Deviations":** Security is not a "set and forget" task; policies must be continuously validated against real-world traffic to ensure that they are functioning as intended. This chapter addresses the use of network telemetry and behavioral analytics to detect anomalies. It covers behavioral modeling, custom security events, and automated response management to quarantine threats the moment they deviate from established baselines.
- **Chapter 10, "Segmentation Maturity Model and Scorecard":** Without a way to measure progress, organizations cannot identify their security gaps or justify further investment in segmentation. This chapter introduces a segmentation maturity model that is based on the Capability Maturity Model (CMM) framework. It provides

detailed scorecards for leadership, architects, and administrators to evaluate their current state and create a roadmap toward a fully optimized environment.

- **Chapter 11, "Reference Architecture":** A robust security strategy requires a multilayered approach where no single failure can compromise the entire organization. This chapter synthesizes the book's technical concepts into a reference architecture. It uses a case study to demonstrate how to layer identity, macrosegmentation, and microsegmentation across campus, WAN, and data center domains for end-to-end protection.
- **Chapter 12, "The Future of Segmentation":** As threats become more automated and move at machine speed, human-managed security will no longer be sufficient. The final chapter of this book addresses the emergence of nanosegmentation (which involves process- and API-level control) and the role of AI-driven autonomous security. It explores the shift toward self-defending networks that can dynamically re-segment themselves in real time to mitigate emerging threats.

Figure Credits

Cover image: Rayhan/Adobe Stock

This page intentionally left blank

Chapter 1

The Segmentation Mindset

From the early days of networking and connectivity across mediums, segmentation has been a constant. Much as with plumbing, there has always been a need to prevent flows of information from interacting with one another, to ensure that they don't contaminate or interfere with one another. Methods to introduce segmentation and separation have varied and evolved over time, typically based on the use cases and needs of those designing the flows.

As segmentation has evolved, it has developed into types such as "macrosegmentation," or the use of broad control mechanisms to group endpoints into categories or logically separated entities, and "microsegmentation." Microsegmentation has been a significant challenge for most organizations to overcome, mostly because of its variable definition for a user, department, or organization. In most contexts, parties can agree that microsegmentation describes the concept of preventing two endpoints from interacting with one another except in explicitly defined cases. The two endpoints could be within a campus, a building, or a floor, or they could even be adjacent to one another and connected to a switch.

As organizations have struggled to apply consistent and effective segmentation—particularly at a granular level—new frameworks have been needed to address the complexity and dynamic nature of modern networks. For example, the Zero Trust model provides a structured approach for organizations to systematically address segmentation and reduce risk across diverse environments.

In their book *Zero Trust Architecture*, Cindy Green-Ortiz and her colleagues describe a five-pillar model for approaching the concept of Zero Trust in a network architecture, where the pillars are policy and governance, identity, vulnerability management, enforcement, and analytics. Many of the aspects of these five pillars are conceptual in nature, and the enforcement pillar is the focus for most organizations. It is common for readers of *Zero Trust Architecture* to start working on changing the organizational mindset to ensure that they have policies and governance, centralized logging of decisions made on the network, and endpoint agents to detect and resolve vulnerabilities. However, network

administrators usually face unique aspects of organizations and inherit brownfield networks—that is, networks already built out in architecture and common practices. They are left wondering how to construct an enforcement approach that works best for their organization and its unique concerns. The enforcement mechanisms typically must take into account macrosegmentation and microsegmentation and the unique challenges associated with each.

This book builds on the Zero Trust foundation focusing on enforcement in real-world conditions. It offers practical strategies and a measurement framework to help organizations assess and evolve their segmentation journey. The following chapters discuss the unique aspects of organizations that require different enforcement techniques, applications of technology, and methodologies. This book also offers a measurement technique for evaluating the segmentation approach and mechanisms used by an organization to help plan or evaluate what is implemented and whether it aligns with the organization's goals.

Pillars of Zero Trust

As we already mentioned, *Zero Trust Architecture* describes five pillars, or required capabilities, for an organization starting a journey with Zero Trust. Figure 1-1 illustrates these pillars, and subsequent sections cover them in turn to provide an overview.



Figure 1-1 *The Five Pillars of Zero Trust*

Policy and Governance

It is important to recognize that strong policy and governance are foundational to both macrosegmentation and microsegmentation. Clear policies determine how broad groups (macrosegmentation) and individual endpoints or applications (microsegmentation) are

separated, governed, and controlled throughout an organization. Ultimately, effective policies and governance enable organizations to define and enforce segmentation boundaries at both the macro and micro levels, and they also ensure that these practices are sustainable and aligned with business objectives as the organization evolves. Figure 1-2 shows the various stakeholders in the success of segmentation.

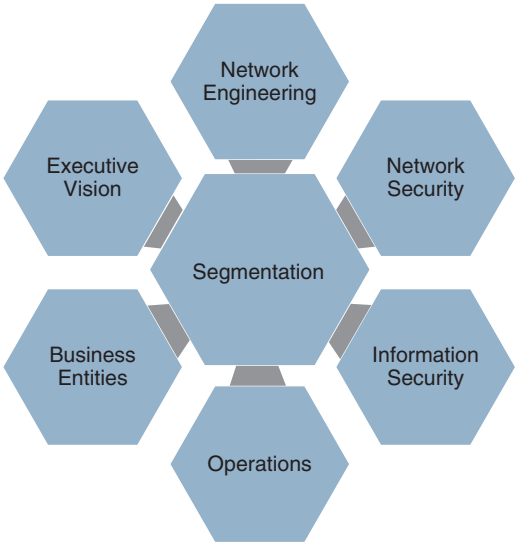


Figure 1-2 *Stakeholders in the Success of Segmentation*

The most challenging aspect of implementing segmentation, as organizations soon find out when embarking down the strategic pathway, is not the technology used to segment a network. It is what engineers commonly refer to as "Layers 8, 9, and 10." These layers represent politics, money, and culture as they relate to the design of a network:

- **Politics:** As it relates to segmentation, the politics of an organization—or just the shared preferences on how technology "should" work—can be extremely challenging to work through. Environments such as manufacturing, healthcare, and finance are at the forefront of this challenge. In any organization where the bulk of the profitability comes from departments previously given the most political clout, changing the mindset to encourage better security practices is the first major mountain to climb on the journey toward segmentation. Political clout brings challenges. For example, in a department that has a lot of clout, users may be allowed to bring their own devices, as long as the devices have a minimal amount of protection running on them. Often, such a level of protection does little to protect the organization; it protects only the endpoint—and typically very poorly. We describe political challenges later in the book. For now, it is important to see that overcoming the political environment that exists takes time, patience, and buy-in from executive leaders to change the shape of how business and technology accommodate each other.

- **Money:** Any department that has attempted to purchase state-of-the-art hardware, software, or processes knows that battles over budgeting for a new addition to a network can be challenging. Often, other departments see that a new capability exists and want to use it but don't want to share the budgetary burden for its use. Staff might need to be added to support new features and capabilities on a network, and that costs money. Network administrators can expect pushback when additional overhead is added to gaining access to areas of the network that were previously protected with user credentials only. Those within an organization—including network engineering, network security, information security, operations, and business owners—need to work in concert across departments to ensure that as new segmentation technologies are implemented, each department understands its role, in terms of supporting the new technology and also the financial burden. As part of the segmentation journey, it is important to account for staff, process changes, time to work through processes, and support for audits of the segmentation methodologies implemented.
- **Culture:** Most organizations are deeply attached to "how we've always done things." In discussions about changing the current approach to properly implement segmentation, the loudest voices are typically those of the individuals who dissent. Teams are overburdened, workloads are higher than they have ever been, and organizations are growing exponentially; those who have the most tribal knowledge are likely to push back when an organization adds a new technology to learn or a new process to work through, or when it needs to plan for how to adopt new approaches to solving problems. Even so, especially at the beginning of a segmentation journey, an organization should entertain new approaches to protecting itself, even if only done in small areas of the network at first before spreading to other areas with greater impact. This could be thought of as a segmentation "liberation," where the organization explores a new approach in order to determine whether it can overcome the lessons learned from the approach, whether the burden is too large for the business to accommodate, or even whether the cultural climate of the organization will prevent the approach from being successful.

The policy and governance pillar of Zero Trust sets the stage for the technological approaches used to segment a network and may be the hardest to overcome in many organizations. After all, the ability to implement policy in any organization could be thought of as being downstream from the organization's culture. Significant buy-in is required of the organization's leadership to ensure that policy will be adhered to, even for the most senior executives, who will need first-level helpdesk assistance to get new endpoints onto the network.

The organization's security and data policy can significantly impact the cultural landscape if that landscape is malleable and influenced from the top down. Senior leadership, both managerial and technical, can enhance adoption by emphasizing the necessity of segmentation as a mandatory practice. Beyond the security and data policy, it is crucial to establish standards, baselines, and guidelines that detail the "how-to" of implementing segmentation. It is important to ensure that segmentation is ingrained in the culture and becomes an integral lifecycle stage for all organizational activities. *Zero Trust*

Architecture provides guidance on the aspects of policy and governance to focus on when implementing segmentation, and you are encouraged to look into those aspects.

The subsequent text brings the five pillars to life, describing them in more detail, and Figure 1-3 shows how the five pillars work in a never-ending cycle throughout the segmentation journey.



Figure 1-3 *Pillars of Zero Trust in a Never-Ending Cycle*

Identity

The first major technological focus area when beginning a segmentation journey is the attribution of identity to endpoints throughout the network. The classic approach to the identity pillar has been to associate an active directory or LDAP object to every endpoint on the network. Organizations quickly find themselves overwhelmed with identity creation tasks as the few admins with LDAP access are overloaded with onboarding tasks. Requests for every new endpoint that needs access pile onto discovering what already had access. Endpoints with existing access also need to be onboarded into the identity store. This is further complicated by political pressures to get endpoints onto the network quickly and with little overhead, bypassing processes meant to understand whether endpoints are expected and allowed on the network.

Establishing accurate and granular identity is essential for both macrosegmentation (grouping entities based on shared identity attributes) and microsegmentation (where policies are enforced at the level of individual users, devices, or workloads). A robust identity framework ensures that segmentation controls are tailored and dynamic, supporting both broad and fine-grained security outcomes. In summary, a strong identity pillar supports the effective grouping and management of entities for macrosegmentation, and it also enables precise, context-driven microsegmentation, ensuring that access and

controls are tightly aligned to the unique attributes and behaviors of every user, device, and workload on the network.

Identifying endpoints, devices, and workloads—which we refer to as *entities* from here on—on a network is as much a process as policy and governance drafting is, and it involves very similar challenges. As discussed earlier in this chapter, in the section "Policy and Governance," every entity should be accounted for in the identification phase, which relies on the same buy-in and involves the same political needs as policy and governance.

For each entity on a network, the organization should take into account major aspects and track them in an attribution engine of some sort. In recent years, this has been done in ServiceNow or similar asset management databases, which allow for tracking of these attributes and can also serve as centralized databases for identity technologies, such as Cisco Identity Services Engine (ISE), to pull information from and use in policy decisions. These are some of the attributes that are commonly tracked for an entity on a network:

- **Entity owner:** The entity owner is the user or department that purchased the entity or that has been assigned to access the entity or use it to accomplish a business outcome. The entity owner may be different from the entity manager. For example, a device may be a personal device such as a BYOD (bring-your-own-device) endpoint that is purchased, managed, updated, and maintained by the same user. On the other hand, the owner may be the end user of a PC that is provided by the organization, where the manager dictates purchasing, updates, and ultimate retirement of the device. The owner attribute becomes extremely important in times of incident response, when it is crucial to know who is expected to be using the device; the organization needs to be able to ensure that the owner can be contacted and given next steps when the entity has a policy applied to it to prevent further risk to the department or organization.
- **Entity manager:** The entity manager, as we previously mentioned, is the person, department, or organization ultimately responsible for ensuring that the endpoint conforms to the policy and governance set out for the entity. When an entity is found to be out of compliance with policies and governance, it typically falls on the entity manager to remediate the noncompliance to allow the entity to perform tasks required to fulfill its business justification for being on the network. In a corporate environment, much of the entity management is performed by a desktop engineering department as part of the operational role described earlier for a successful segmentation approach. However, in an organization such as a large medical provider with research endpoints on its network, a manufacturing organization that allows for third-party vendor-programmable logic circuits, or an energy provider that often leaves sensors unattended and autonomous for years at a time, the management of devices may be distributed to individual departments or even individuals, making for a much more complex management approach.
- **Entity support team:** Who gets the first call when an entity is not doing what it is supposed to do? The support team. This might be a helpdesk that a corporate user contacts to open a ticket for support when their device is not behaving as

expected. The responsible support team—or perhaps unlucky individual—should be documented for two reasons:

- To ensure that all teams are knowledgeable about a change to the access provided to an endpoint when it is not behaving as expected
- To ensure that the proper department or individual can be contacted when the device does not behave as expected and to ensure that standard approaches are used for support

A standard approach is to maintain common knowledge databases of behaviors that can be easily updated based on past experiences. There may be a billing aspect involving who is billed and who receives the funds when a need for support arises.

- **Expected lifecycle of the entity:** Entities do not, and should not be expected to, have an unlimited lifespan on a network, despite how legacy corporate networks handled them. A modern car has a lot of features that weren't available in a classic Chevy Bel Air, including seatbelts, built-in GPS, real-time support subscriptions, and a blind spot monitoring system. Similarly, modern networks have evolved to need new technologies, protocols that support new business outcomes, and the ability to process more data at a quicker pace. As networks evolve, it is necessary to sunset support for technology over time, retire the technology, and upgrade to more modern entities to perform the tasks. It is important to estimate the lifecycle of an entity at the time it is placed on the network in order to plan for the entity's future (including budgeting) and ensure successful retirement and replacement of the entity. This lifecycle should be considered as part of policy decisions, as we will discuss later in this book.
- **Identifying attributes of the entity:** Attributes of entities need to be identified. For example, a standard PC may have multiple identifying attributes, including its MAC address, hostname, make, and model. The most consistent attribute—in this case, the MAC address of the endpoint—should be stored in the asset management database. Technologists reading the preceding sentence will undoubtedly scream "But I can use a utility to change my MAC address to whatever I want on the network!" However, consider that the MAC address is a consistent attribute that should be unique across multiple devices. The identifying mechanisms and processes used should include duplication detection to ensure that duplicates are expected or eliminated based on a device's expected behaviors.

  The MAC address should be considered a singular attribute within what has been dubbed a "contextual identity" of an entity. Further consideration for what a device "looks like" when on the network—including whether a user is seen as using the endpoint, behaviors of the endpoint that may indicate its identity (such as HTTP headers or DHCP attributes), the medium via which it joined the network, and the location it joined from—should also be used to develop context around the identity of an endpoint. Additional factors, such as how an entity behaves when it accesses the network and whether it requires additional dependencies not recorded, will dictate the access the endpoint gets in future phases of the enforcement approaches described later in this book.

- **Known dependencies:** Documentation of what port, protocol, or entity interactions a given endpoint requires to perform its business functions is rarely exhaustive, and many times it is inaccurate beyond the basic functionalities of the endpoint. Often, tribal knowledge—that is, knowledge gained by operationalizing an entity and then documenting it to make the supporting team's or individual's life easier in future—is far more accurate. Therefore, especially for incident response, it is a good idea to fill in a basic free-form text field to indicate how an endpoint interacts with other services on the network, which protocols it utilizes, and which shared services it relies on when evaluating the endpoint's asset management database entry.
- **Last evaluation date:** Technical debt is a problem for most organizations. As part of an asset management database, the last evaluation date for each entity should be documented. An evaluation of an entity can be as simple as the last time the entity checked into an endpoint management system, such as Microsoft's System Center Configuration Manager, or it could be the last time a medical device was taken off the floor to have its firmware updated. The last evaluation date should not impact the expected lifecycle of the device. The lifecycle should continue to influence budgetary decisions and management's value in replacing hardware by tracking how long a device has served its purpose; the last evaluation date, on the other hand, should merely serve as a check that a device was updated to mitigate potential vulnerabilities in the time period since the entity was introduced or last evaluated.

Vulnerability Management

Vulnerability management is typically implemented in the form of endpoint software or an active means of evaluating what risk an entity poses to the network based on the access provided to it. In many product sets and for many organizations, this may be referred to as the "posture" of an endpoint. However, as many organizations have pointed out in the time since *Zero Trust Architecture* was published, vulnerability management has as much to do with the enforcement mechanism applied to an entity as it has to do with the identifying criteria associated with the entity.

Effective vulnerability management informs segmentation decisions at both the macro level—by identifying and grouping high- or low-risk entities—and the micro level, where dynamic policy adjustments can be made in real time, based on the posture or risk score of each individual endpoint, user, or workload. In summary, vulnerability management is about identifying risks, and it is also about enabling adaptive segmentation—ensuring that both broad group policies and granular, per-entity controls are informed by the real-time risk posture of every asset on the network.

Vulnerability management should be thought of as evaluation criteria for determining the potential risk of an entity connecting to the network. It should include measurement of risk, and a score, based on what the entity can provide to the network, for how sure the organization is that the endpoint is what it says it is. In the Zero Trust model, this risk score isn't just a passive measurement; it is actively used to drive enforcement decisions. Instead of relying solely on static, role-based privileges, Zero Trust architectures

use continuous assessment of identity, endpoint posture, location, and behavior to adapt access levels in real time. As the risk score for an entity changes—whether due to new vulnerabilities, suspicious behavior, or changes in location—enforcement mechanisms can dynamically raise or reduce that entity's access to ensure that segmentation policies remain tightly aligned with the organization's current risk landscape. This evaluation will be a large aspect of this book, but it basically involves the following questions:

- **What information can an organization gather from the endpoint to evaluate its identity?** The vast majority of endpoints connected to networks today have no centralized management, and few have the ability to actively identify themselves to the network. Those which can identify themselves may be using an easily changeable mechanism such as a hostname as opposed to verified credentials. In these cases, workarounds such as MAC Authentication Bypass (MAB) are employed, but are easily spoofable, as is hostname-to-DNS mapping.
- **How does this endpoint interact with common services?** An identity database should include a measure or baseline of expected interactions the endpoint uses to fulfill its business function. The port and protocol baseline the endpoint uses to interact with management systems and peers alike should be evaluated, if not mapped, to ensure that in the case of any significant deviation from this standard, the endpoint can be automatically or manually limited in its access. Chapter 5, "Microsegmentation," describes how to identify baseline interactions for specific types of entities.
- **What is the nature of the hardware this entity represents?** The average PC will of course have multiple applications running on it, all with unique needs to access management systems, peers for data exchange, and services from which the PC consumes data for processing. The growth of Docker containers, the availability of commercially free virtualization hypervisors, and use of virtual environments by programming languages such as Python were previously reserved for data center hypervisors running on an array of hardware. Today, they are considerations that influence what vulnerabilities an entity is subject to.

Enforcement

Enforcement is the "meat and bones" that helps determine how much segmentation a network requires, how to evaluate whether the organization and network are ready for segmentation methodologies implemented based on their impact for operations, and where to apply the best level of segmentation across multiple use cases. Throughout the enforcement evaluation and application process, it is important to remember that there is no "silver bullet" or even single place where an enforcement mechanism can be applied to prevent an entity from posing risk to a network. As many organizations have found out through multiple false starts at enforcement, planning is key to successful application of enforcement, and successful planning begins with understanding the need for segmentation, the possibly long timescales involved in the segmentation journey, and the level of impact that a particular network can tolerate.

Enforcement is a battlefield that consists of multiple layers of defense. Each enforcement control point is a test of whether the strategy, planning, and policy constructed in theory can survive the onslaught of real-world usage. The enforcement points are layered from the access layer with 802.1x and SGTs to the distribution layer with ACLs and SGACLs; in the data center through microsegmentation; in the public cloud with security groups; with cloud-native infrastructures using service meshes; and for remote access using SASE/SSE. Each of these layers demands its own planning, stakeholders, strategy, policy, measurement, and evaluation of segmentation. What works in one layer may not necessarily work in another, but consistency is critical. The following chapters cover these enforcement techniques and layers in more detail, culminating in the reference architecture provided in in Chapter 11, "Reference Architecture."

Enforcement is the practical pillar that brings segmentation to life: It translates high-level security policy into concrete controls, enabling both macrosegmentation—by creating broad boundaries between groups of entities—and microsegmentation, where fine-grained rules are enforced at the level of individual users, endpoints, and applications. Ultimately, effective enforcement ensures that segmentation strategies—whether broad or highly granular—are consistently applied, dynamically adjusted, and aligned to the unique risk profile and operational requirements of the organization.

Analytics

Perhaps the second-most-important planning step is analytics. Analytics will be defined based on organizational goals for segmentation and will influence every other step of the segmentation journey. Analytics will be the reference point for future iterations, which will continually evolve as segmentation expands and is redefined for an organization. To clarify, analytics is the understanding of the enforcement mechanism applied to an entity as it came onto the network, why this enforcement result occurred, the conditions evaluated to reach the result, and the subsequent impact on both the risk posed to the network and ability to accomplish its business purpose. Analytics is the lessons learned from ongoing evaluation of each of the five pillars. As a result, this analysis is done after each pillar is evaluated to ensure that the next time the pillar is evaluated, lessons learned can be gleaned and applied to every iteration thereafter to ease the burden on the user, administrator, manager, or support personnel for the entity in question.

The key performance indicators (KPIs) that are analyzed will differ from organization to organization. Analysis will be based on the goals the organization has for its segmentation implementation. Chapter 9, "Validating Policies, Monitoring Enforcement, and Responding to Deviations," outlines a sample approach to evaluating the success of segmentation. As an initial success metric, analysis should evaluate whether users and endpoints have access to key resources required to perform actions related to their roles, based on understanding of endpoint interaction requirements. This access, while also limited at the application level via mechanisms such as role-based access control, should be limited by the network to ensure that key resources remain accessible. Resource limitations should be applied in a phased manner, increasing over time or as use cases call for such needs. For example, the initial phases would involve permitting full access, and later stage mature phases would feature granular access at the port and protocol levels.

Robust analytics enable organizations to assess the effectiveness of both macrosegmentation and microsegmentation strategies by providing insights into how segmentation controls perform across large groups and individual entities and by informing continuous improvements to segmentation policies and enforcement. By leveraging analytics, an organization can ensure that segmentation remains adaptive, evidence based, and aligned with both security and business objectives at every level of the network.

Beyond the Five Pillars

While the five pillars provide a structure and great starting point for what organizations need to consider as they work through their Zero Trust journey, significantly more guidance can be posited from the architectural perspective of how to plan and build for a properly segmented network. After all, most organizations have been running their networks for decades, making incremental improvements as they encounter new technologies, regulatory needs, and threats. Organizations use various techniques to apply differentiated access to entities on the network, including VLANs, firewall termination, intrusion prevention systems, and, in more advanced networks, endpoint identity.

However, technological advance should not be measured by how segmented a network is or how many autonomous systems entities belong to. Many organizations now overuse segmentation, separating each system of interacting entities into a virtual routing and forwarding (VRF) instance, for example. In such cases, planning for a repeatable process creates such a complex structure that it can only be supported with high-end hardware, significant operational resources, and enough knowledge by all those who support the architecture that the time to resolution for architectural changes or daily operational challenges increases well beyond what the organization can handle to remain nimble.

An organization should strive to implement segmentation methods that fit the organization and the challenges it plans to address; these methods need to scale with the organization as new use cases, entities, or regulations are applied and require resolution. This book addresses how to use the macrosegmentation common to most networks and overlay it with microsegmentation. The combination of these two strategies makes it possible to segment a network in such a way that business outcomes can be accomplished; it also protects the network when one method of segmentation may not be able to be applied due to common challenges or troubleshooting techniques. Moreover, microsegmentation provides a protective cover that prevents lateral movement of malicious actors even if they are within the network; it does so by making the attack surface smaller and building in a “deny all and allow only what is trusted” approach.

Addressing the Problem Head On: Why Do Segmentation Strategies Fail?

A large healthcare organization needs to validate to its auditors that its segmentation strategy prevents common attacks. Its primary motivations for creating a segmentation strategy are to meet strict regulatory compliance requirements and protect sensitive

patient data by reducing the risk of malware spread and other lateral attacks within the environment. The initiative is not only a technical endeavor but also a direct response to external pressures from auditors and increased awareness of industry-wide cyberattacks targeting healthcare systems. With criteria in hand, showing its ability to prevent malware spread between endpoints, the healthcare organization attested in writing that each endpoint on the networks within its largest hospitals would be subject to active authentication as it connected. Active management devices, such as corporate endpoints, were managed through Microsoft System Center Configuration Manager (SCCM) and configured via group policy to present the user's credentials to the network. A custom attribute was used as part of DHCP address allocation to indicate that an endpoint was a corporate PC and profile, or fingerprint, the device accordingly. Medical devices on the network were onboarded into an LDAP database, in which MAC addresses were populated as usernames and were used to authenticate the endpoint. Each device was allocated to a VLAN dynamically, based on its device fingerprint, and each device belonged to a VRF instance, adding to the segmentation approach. Downloadable access control lists were applied to endpoints to give them access to compute resources locally and data centers remotely. The customer was actively working toward building a micro-segmentation strategy to prevent intra-VLAN communication. However, the organization still failed a third-party audit. What went wrong?

As the healthcare organization began to implement its segmentation strategy, it did so on a floor-by-floor basis, with an approach that seemed gradual but that affected multiple departments and critical uses at once. Operations teams received an influx of calls from physicians and researchers who were unable to get their endpoints onto the network. Despite the network engineering team's best efforts to ensure that each endpoint was allowed basic access to the network, even if active authentication failed, the team didn't plan on endpoints running software that had little fault tolerance. The software running on these endpoints never planned to be on a network that would deny the hardware endpoints immediate access to resources once connected. For example, imaging software immediately reacted to changes in network state. When the network adapter came up on the wired medium, the software would immediately try to connect to the server that the client interacted with. The software would indicate that the endpoint was disconnected from the network during even the slightest delay in gaining connectivity to its server. The user interpreted this delay as not having network access and attempted to reconnect to the network in a perpetual, problematic cycle.

Because the imaging software was extremely large in size and required significant processing power, the endpoints were purchased from outside the standard catalog to allow for customization in hardware specifications, with higher CPU and greater amounts of storage and memory. These endpoints were therefore imaged not with a standard corporate image but with a customized image that provided administrative access to the end user. Many of the users saw active authentication as a burden and overhead they did not want to deal with, and an unfounded fear that the active authentication would cause a delay in getting onto the network began to spread across many of the more technical users. These users turned off active authentication either proactively or reactively, without evidence of system impact.

As word spread that the capability to turn off the active authentication was a possibility, more users did so as they perceived delays in getting onto the network. Passive authentication required a longer delay to validate credentials, causing anxiety in users who interpreted software error messages as not having connectivity. As they unwittingly made the authentication process slower by transitioning the endpoint from 802.1x authentication to MAC Authentication Bypass, helpdesk engineers found an easy solution: Remove authentication commands from the switchport to which the endpoint connected and assign a static VLAN address. While this was meant to be a temporary fix to quell the onslaught of tickets from users complaining about software incompatibilities or delays getting onto the network, the temporary solution lacked documentation or technical review and became part of standard procedures ("tribal knowledge"). Out of 4500 ports throughout the organization's largest hospital, 2000 of them now lacked authentication. The audit uncovered this failure.

Much like in a climb, where having the right equipment matters and strategy is crucial, the success of segmentation depends on technology and also on the approach taken. Had the organization begun with a "low-impact mode" segmentation strategy—gradually introducing controls for visibility and feedback rather than jumping directly to full enforcement—it could have identified user resistance and technical issues early on and made course corrections before widespread noncompliance became ingrained.

This scenario underscores several key lessons:

- Introducing segmentation in phases enables IT teams to gather real-world feedback and adapt policies before enforcing them organization-wide, reducing the risk of operational disruption.
- Actively engaging users and support staff during the initial rollout helps uncover usability concerns or misunderstandings that could undermine the technical solution.
- Maintaining clear documentation and continuous communication prevents tribal knowledge and ad hoc fixes from becoming entrenched, preserving the security intent of the original segmentation strategy.

Ultimately, a successful segmentation initiative requires not only robust technology but also a thoughtful, flexible rollout plan that balances security, usability, and operational realities across the organization. In addition, any segmentation initiative requires four organizational considerations that will often dictate a segmentation initiative's success: a well-aligned team, buy-in from business entities, executive-driven vision, and alignment of strategy and tactics.

The Importance of the Team

For the major healthcare organization in the preceding example, everything seemed to be correct and on track:

- Endpoints were centrally managed—except for the endpoints that needed more horsepower.

- Endpoints were generally received from a distribution facility with a standard image—except for special use cases that were purchased from outside of the catalog.
- Users were informed of the benefits of active authentication. However, human nature will naturally push back on a removal of resources previously allowed, such as unfettered access.
- The operations team was informed of the active authentication and benefits of having it in place. However, when overwhelmed, the helpdesk fell back to legacy networking techniques to solve a challenge.

The most significant weakness for the major healthcare organization was a breakdown in the team dynamic in the face of challenges. As illustrated previously, the five Zero Trust pillars—policy and governance, identity, vulnerability management, enforcement, and analytics—all need to work in tandem to ensure the success of segmentation strategies. Each pillar pertains to specific aspects of the connectivity of the user, endpoint, or application, which needs to be considered. These considerations have a downstream effect on successful implementation of the segmentation strategy. Teams throughout the organization must therefore be aligned. Key to this alignment is an organization-wide understanding of which mechanisms were used to enable each pillar, and why:

- Not understanding the identity of an endpoint prevents knowledge of where and why the endpoint is connected.
- This undermines the ability to minimize the risk it poses to the network.
- This altogether eliminates any ability to enforce a policy against it.

This breakdown is what ultimately led to the failure of the audit. The following paragraphs explore the role of each critical team and underscore their need for alignment.

Information Security

Policies must allow for enforcement mechanisms to be applied to entities on the network. These policies are generally owned by information security, driven by the needs of the business, and are always signed off on by executives who, along with their teams, are governed by the policies. Policies must consider the potential impact of applying enforcement mechanisms to entities and the potential impact of not applying enforcement mechanisms as a risk statement. From a technical perspective, information security must be involved in drafting and approving policies for what it means for an endpoint to be segmented while on the network. This will typically take the form of what requirements or conditions may be needed for an entity to be connected to the network, and the resulting access if these conditions are met. However, these policies must be feasible to the operational goals and needs of the business. Policies that are overly complex, unclear, or vague will be overcome through technicalities. The use of phrases like “may” when applied to endpoint access can mean that specific business owners will push back on why the policies apply to them. In the face of pushback, operations teams or even executives may choose to allow exceptions to multiply. The larger the exception list, the more insecure the organization becomes.

Because the major healthcare organization's users saw active authentication as a negative, they were not in line with policy but still had access to the network. Because they also used machines that were purchased outside the normal purchasing process, there was no oversight or management into whether the endpoints were in line with information security policy. In an ideal scenario, there is a policy, and all business entities that it affects understand it. This eases conversations with business entities when endpoints that do not perform active authentication are prevented from accessing the network, as they should have been in this scenario. The same goes for when endpoint owners are assigned tasks to remediate this misconfiguration.

Network Engineering/Desktop Engineering

Network engineering and desktop engineering provide the first technological means to evaluate the policy decision point and validate whether the endpoint belongs on the network. The endpoint connects to a network access device, and the network access device then prompts the endpoint for credentials to validate its identity. This occurs across wired, wireless, and remote access mediums. At this point, there are three potential identities the entity can provide to the network: a user credential, a machine credential, or no credential at all. Whether the entity provides a user credential or machine credential is a policy decision and technological decision that needs to be made based on the policies required by information security. A managed endpoint running Windows 10 or above can provide a combination of both user and machine credentials or just one or the other. This is commonly used when the policy decision point, such as Cisco Identity Services Engine, is integrated into Microsoft Active Directory or Microsoft Azure, where policies can be pushed to the entity, configuring it to present its credentials to the network from a native piece of software called a supplicant.

The real challenge begins where endpoints, such as the imaging endpoints for the major healthcare organization mentioned earlier, are not managed by a singular entity. It continues when the configurations applied to a port to which that endpoint connects are not monitored or tracked for compliance with a "golden," or known good, configuration. As a best practice, any endpoint that needs access to a corporate managed network should be accounted for through some form of centralized management. To reference that centralized management system, the network access device—switch, wireless controller, or VPN aggregation point—must be configured to reference the authentication against a network access controller, such as Cisco ISE. With advances in endpoint and network configuration management, there is a ubiquitous ability to push configurations to both endpoints and the devices through which they connect to the network. Innovations in being able to push configurations and policies even without a fully managed endpoint has added significant flexibility for network engineering and desktop engineering teams to ensure that policy is always complied with.

Network Security

Unfortunately, blame is often cast on network security and networking teams that are implementing the vision of executive leadership in alignment with information security policies. Even after consulting with the business entities that a technological application

will affect, networking teams are often blamed for implementations. When a problem arises, those in a typical organization often cast a string of blame in the following order:

- "It's a DNS problem. There was a change on the DNS server recently."
- "It's the firewall's fault. It's blocking my access."
- "It's the switch's fault. I can't authenticate."
- "They made a change to my computer, and now it doesn't work as it's supposed to."
- "My computer was never connected to a network."

A network security implementation team must be prepared to test every potential possibility and use case throughout the process of implementing segmentation for the organization's network. This is, by any measurement, a massive task that needs to be broken down into the smallest pieces possible. For common use cases, it is best practice to have a small-scale network that represents major use cases and includes spare endpoints that can be justified in budget as endpoints to minimize mean time to resolution by replicating the most common setup and use cases. In the major medical provider's case, a lack of understanding of the use case and how critical it was to the business from both the network engineering and network security teams was an opportunity for improvement. It could have been tested in a nonproduction environment as a critical use case that would impact implementations and policies if not accounted for.

One common limitation for many organizations is the budget justification for having a nonproduction environment in which to test use cases and technical application of policies. A measurement of mean time to resolution (MTTR) can be used to justify such an investment, showing that testing the use cases well ahead of implementation can ease outages and the need to resolve over long periods of time, as well as the rework that occurs when the solution is removed and readded when time to troubleshoot is allocated. As a tradeoff to this budget justification, these endpoints then need to be used as a test bed for the implementation of new features, even before the pilot that implements these new enforcement abilities in a gradual manner across endpoints. By testing well ahead of time and taking note of how entities—both hardware and software—impact users and manifest their results, operations teams can learn what to expect and how to address challenges that arise while remaining within policy.

Operations

One of the most effective approaches to applying segmentation policies is to think of the network from the perspective of an operations team. Operations is often caught in the middle of the efforts of the rest of the required stakeholders due to its role as the team that gets an entity connected to the network to fulfill a business purpose and has to deal with users who are unable to get their entities connected to the network. Its goals differ from those of other teams:

- The information security team wants to ensure that an endpoint abides by policy and keep it from accessing the network if it doesn't abide by policy.

- The network security team wants to ensure that endpoints that pose undue risk to the network are prevented from accessing the network or at least ensure that their access is limited to only the granular subset of resources they require.
- The network engineering team wants to verify that an endpoint connected to the network has access to the resources it needs to fulfill its purpose on the network.
- The user and application team want to ensure that the entity can process data and information to generate revenue for the organization.

Operations, in this scenario, is arbiter of the endpoint getting onto the network, based on the requirements provided to it from the other teams. Among other challenges, an inability to operationalize their segmentation strategy ended up being a major contributor to the major healthcare organization's failure of its audit.

Cross-Team Collaboration

Ultimately, the operations teams fell back to using a legacy approach to segmentation and statically applied VLANs to endpoint ports. The network security team expected dynamic enforcement mechanisms to be applied to endpoint sessions in accordance with policy. The network engineering team expected configurations to remain consistent after applied. The information security team expected policies to allow for connectivity for the business entity.

In an ideal approach, multiple improvements would have been made during the rollout of segmentation at the major healthcare company. A daily meeting at the time of enforcement, decreasing in frequency over time, would have brought all of the major stakeholder teams together to work through challenges heard and observed. This might have changed policy application or driven the need for new policies relating to the imaging workstations. The organization would have established policies that provided alternative mechanisms that allowed users and endpoints to connect and reevaluate the policies over time. Where the alternative connection mechanism was not possible, an exception process to track and remediate the issue would have been used to close these gaps. A process to truly evaluate the need for exceptions, based on technological feasibility, would have been used to ensure that the number of exceptions didn't grow as long as the policies themselves.

Cross-team communication at the leadership levels is the ultimate success criteria in this scenario. Without this communication, the adaptability of teams to react as they saw challenges arise was prevented. In the scenario, an ongoing conversation would have helped the operations team understand the use case for the endpoint, the measurements of success for the endpoint while on the network, and the restrictions applied to the endpoint and its ability to connect. The downstream implication and need for the endpoint and its use case could then be socialized among major stakeholders. If the owner of the endpoint and use case was unable to effectively describe the use case, or if the security teams were unable to describe the need for restrictions to the business, either could be hard pressed to continue justifying their actions. The operations team could also provide feedback saying there were limited abilities to differentiate the imaging machine use

case from a normal PC. Finally, the teams could discuss how to prevent technicians from autonomously applying a simpler approach to granting the endpoint access to the network, as human nature would be to pursue the easiest possible solution. In this case the easiest possible solution was to provide unfettered access to the endpoint.

A common rule used during the design of segmentation strategies is, "If a college intern—the same college intern who is an operations technician for your network—can't understand your segmentation strategy, the strategy will fail."

Business Entities

The first Zero Trust pillar—policy and governance—will always be the most important. The inalienable truth is that every endpoint and use case must be governed by policy in some way. The reference to this policy is the only way to justify protection mechanisms applied to the endpoint, which ultimately protect the network. It has been said that the most secure endpoint is the one that is unplugged, 9 feet deep in the ocean, with no user to interact with it. As long as users connect their endpoints to the network to perform their business purpose, it is important to plan for how they use the endpoints and assess the risk associated with the endpoints' capabilities, so that a policy can be written to offset that risk.

A major focus of *Zero Trust Architecture* is on forming policies that govern how users utilize an endpoint, identifying the user and the endpoint, understanding the interactions of the endpoint on the network, and applying restrictions to any other interactions. However, the book does not focus on the simple truth that if an endpoint is on a network, someone put it there for a purpose, whether positive or negative. Understanding the purpose of that endpoint being on the network and implications behind that purpose must be the highest priority when forming a segmentation strategy and working within a team to do so.

In the healthcare example, the problematic endpoints connected to the network needed to process imaging data. The imaging data is massive in size and requires significant throughput. An application processes these images, and it is programmed in such a way that it requires instantaneous access to the network; otherwise, users feel as though the endpoint is disconnected or is taking significantly longer to access the network than before active authentication was introduced.

For the major healthcare organization, understanding how users and departments utilize their machines—either by conducting interviews with key stakeholders or gradually introducing the segmentation plan to these devices rather than all at once—would have helped isolate the problem and minimize its impact. In scenarios like this, it is beneficial to have relationships with key stakeholders who are considered critical to each business entity who may be able to have their endpoints subjected to segmentation approaches in a preliminary phase. A phased rollout like this ensures that there is little to no known impact by the time the rest of the business entity has the same restrictions applied.

In addition, communication to the business entities was significantly lacking in this scenario, as evidenced by the rumors and fearmongering that began when a set of critical

users were affected by the applied restrictions. When the imaging techs' software was affected due to their unique use case, discussion quickly moved away from analyzing the root cause of the issue and toward the change that was made to the endpoints. The affected end users changed the behavior of the endpoints, unknowingly causing a larger issue. Especially in cases like this, communication to the business entities is critical, as is developing relationships with those leading the business entities. In this scenario, the business entity leaders with "an ear to the ground" could have bubbled up the observed issue as soon as impact was encountered so that an analysis could be performed to understand the drivers behind the challenge. Armed with this knowledge, the organization could have made small changes to endpoints or processes to overcome the perceived long connection time.

Executive-Driven Vision

A key weak spot to the major healthcare organization's ability to implement its segmentation strategy was a conflicting priority in the minds of the various teams. Despite many of the other challenges—endpoints purchased outside the catalog, software errors that should have been tested, and configurations being removed from switchports—an overarching theme is a lack of understanding of why the efforts to segment the network were so important. When embarking upon an initiative that can have these demonstrated, widespread impacts on all the stakeholders described here, it is especially important to ensure that a top-down message is used to articulate the mandate that segmentation be successful. Where initiatives that span such a wide space of the business start at a managerial level or even director level, the lack of authority to mandate the initiative based on vision will ultimately cause it to fail.

Policy and governance are only as impactful as the executive vision that drives the policy and binds all strategic leaders to it. Key to the success of a segmentation strategy is having someone with both vision and authority to delegate execution of that vision. In the typical enterprise, this implies a need for buy-in at the CEO, CIO, CTO, or CFO level who can then delegate to leaders within each involved business unit to execute on the vision. The hope for this delegation is that as strategy to accomplish the vision is broken into tactics, the most knowledgeable individuals on the impacts of these tactics can help shape and execute the vision successfully. In addition, the shared reporting structures of those most knowledgeable individuals can be used to mediate and resolve conflicts that may later impact the success of project aspects.

Other Considerations Beyond Standard Teams

It's important to recognize that segmentation strategies and the operational challenges that come with them extend far beyond corporate endpoints. As organizations expand into more complex environments, including data centers, cloud platforms, and AI workloads, macrosegmentation and microsegmentation become even more critical. In a data center, macrosegmentation may be used to group application tiers or environments, while microsegmentation may be used to restrict east–west traffic between workloads. In cloud environments, segmentation must account for dynamic scaling, multi-tenancy, and the

ephemeral nature of virtual resources. AI workloads, which often require access to large volumes of sensitive data and specialized compute resources, demand careful segmentation to protect intellectual property and comply with regulatory mandates. Ensuring that segmentation policies are clear, practical, and operationally feasible in these environments is just as vital as in traditional enterprise networks. As you plan segmentation in increasingly diverse environments—from corporate networks to data centers, cloud environments, and AI workloads—prioritizing clarity and operational simplicity early on will help ensure a smooth journey from macrosegmentation to microsegmentation and support the organization's evolving needs.

Segmentation needs to span multiple technological domains. It is therefore important to start small with segmentation planning, and evolve the segmentation plan as it grows to fit larger use cases and domains. Each of these use cases and domains will have separate teams that focus on them specifically. Strategically, a consideration for what criteria can be used to identify assets (part of the second pillar of the Zero Trust Architecture model) across the organization should be developed considering campus use cases first. This can then expand to other domains, such as the data center and WAN. By doing so, the most prevalent physical endpoint type in the network is used as a defining method for the criteria that can be used to identify assets. From experience, as this identification model expands beyond the campus into the data center, the overlay of applications on top of physical endpoints creates a need to add to the campus definition of "assets," rarely taking away from it. That said, considerations should be made to ensure that policies used to define "identity" are malleable and dynamic enough to allow changes based on the multiple, distinct needs of each environment. This ultimately quells the common ask to "boil the ocean" and try to define identity criteria, groupings, and policies for all environments and use cases at once.

No single model will work for every organization, and the use of the five capabilities of the Zero Trust model, with an especially large focus placed on identity, vulnerability management, and enforcement, must influence what works for organizational needs. The following criteria should still be top of mind during these planning stages:

- Ensure that a governance policy exists
- Identify each workload, user, or endpoint
- Understand the interactions of each workload, user, or endpoint within, across, and outside of its system
- Apply an appropriate enforcement mechanism
- Analyze whether this accomplished the needed segmentation requirements

In addition to politics, culture, and team dynamics, other considerations such as the history of an organization and drivers behind past decisions may also need to be considered as part of the alignment of tactics and strategy. A common approach to networking in decades past was to charge business entities based on the number of ports consumed. This led to a large challenge with "shadow IT," or devices being connected to expand the number of ports allocated to a business entity without the broader operations team being

aware of the devices in question. In many cases, this then caused entirely new operations teams or staff to be responsible for business entity networking, the lifecycle of a new network gateway, configurations downstream from the network gateway to provide connectivity to endpoints used by the business entity, and upstream connectivity for all endpoints relying on this networking device. In addition, a "sworn to secrecy" attitude was adopted by these teams, creating a more challenging support environment as well as varied standards across business entities on how connectivity was established.

Aligning Strategy and Tactics

The focus on segmentation should not focus on a singular approach for any of the required elements—identification, evaluating risk, or which enforcement mechanism might be most prudent. To help illustrate the need for broader thinking, consider that segmentation is not a new concept born from modern technology. Even in the medieval era, segmentation was applied in the design of castles to protect the king and other important assets. The castle-and-moat architecture created layers of defense: outer walls, gates, drawbridges, courtyards, and inner keeps, each serving to delay and deter attackers. This layered physical security approach is essentially replicated in the perimeter security used in data centers today. However, just as a castle's defenses depended on the alignment and cooperation of the entire kingdom—from guards to architects to the ruling council—modern segmentation strategies must be driven by coordinated vision and buy-in across all organizational levels.

Segmentation is often viewed as a tactical solution to specific problems. While it addresses immediate security concerns, its true potential is realized when it is integrated into a broad, long-term security strategy. This strategy needs to exist and persist across business units, from operations all the way up through engineering, lines of business, and executive management. A common mistake is taking a "set it and forget it" mentality toward segmentation. There are three main issues with current segmentation techniques:

- **Access control for internal network segments is typically designed within the bounds of how the organization has approached the problem in the past.** It typically follows a tactical approach that struggles to adapt to evolving business needs or is restricted by how the design was originally intended to function.
- **The trust evaluations that form the basis of access policies are often static and quickly become outdated.** By neglecting the ongoing nature of a segmentation strategy and the need to continuously evaluate new use cases introduced to the network, the same policies are inappropriately applied to new use cases. This is similar to trying to use the same five firewall rules to govern access into and out of a network in perpetuity. Inevitably, rules for access will need to change to provide granular access to resources across the varied use cases introduced as new endpoints or users begin to use the network or use the network differently.
- **Access control policies are difficult to enforce effectively due to the absence of advanced (microsegmentation) security components from the data center to the network edge and an inability to efficiently monitor and manage these**

components. As with a firewall used for segmentation, ongoing monitoring and understanding of additional use cases will be required for segmentation to succeed. This includes both onboarding techniques for endpoints and ongoing evaluation of traffic flows.

These challenges frequently arise because management, network engineering, and operations teams design the segmentation architecture without sufficient focus on process and strategic needs to ensure that the network remains protected, as opposed to protecting it at a moment in time. This is magnified by the dynamic and rapidly changing business needs and hybrid environment that have become the norm for modern infrastructure. Chapter 2, "Alignment with Business Outcomes," explores the impact of these three main issues with segmentation techniques.

Maturing in Segmentation Strategy

Teams that support the segmentation goals of an organization rely on each other. Without vision and strategy, there cannot be a successful implementation of tactics and technology. Organizations must position teams to drive the success of their segmentation strategy in order to measure and evaluate their segmentation strategies. However, vision should also focus on starting small and evolving the strategy over time, as an organization matures in operations and technical capabilities.

One common limitation that is put on organizations when designing a segmentation strategy is a problem statement such as, "The network should be broken into no more than seven use cases." In most organizations, common segments that are observed would be some or all of the following:

- Corporate endpoints (user endpoints, printers, IP phones)
- Building management systems (HVAC, thermostats, pumps, elevator control devices, irrigation systems, parking garage sensors)
- Security endpoints (IP cameras, badge readers, emergency call boxes, PA systems, gate control systems, sensors)
- Lab devices and guest/demo endpoints
- Business entity–specific devices (medical endpoints, manufacturing endpoints, trader turrets, retail endpoints, point-of-sale systems)
- Business entity–specific servers
- Quarantined endpoints

When creating a segmentation strategy, this common breakdown will help organize endpoints into use cases that can have common policies applied to them at a high level. However, it is also common to get pushback on such a limitation and questions why many technical consultants see it as the best practice. The idea that the organization needs far more granular segmentation policies applied to endpoints to prevent unauthorized access

does not take into account the ability to operationalize the implementation when it is rolled out throughout the organization. The challenge with the mindset that the segmentation strategy needs to be extensive to fit the organizational goal may be correct—but not until a much later point in the segmentation journey. By approaching a segmentation journey with the granular segmentation mindset from the beginning, an organization typically cannot describe its segmentation policies to a college intern and will therefore overwhelm its operations team—setting up the team for failure from the very beginning.

The advancement of modern security frameworks—particularly Zero Trust—has been a catalyst for many organizations in increasing their segmentation maturity. These frameworks encourage, and often require, organizations to coordinate across teams and functions to achieve more granular and adaptive segmentation. This push is driven by evolving business requirements and the need to support modern workloads and also by growing compliance obligations in industries where data privacy and protection are paramount.

Later in the book, we present a method to evaluate an organization's segmentation strategy. Evaluation is critical, but attempting to evaluate a segmentation strategy before gaining buy-in and planning how teams will collaborate to make segmentation successful will inevitably lead to a much more challenging and lengthy approach to properly segmenting the network. By leveraging the principles of frameworks like Zero Trust, an organization can more effectively align team efforts, streamline its segmentation journey, and adapt to future security and compliance demands.

A Look Ahead

Most chapters in this book build on the previous chapters; however, some chapters stand alone in their recommendations and the topics they cover.

Chapter 2, "Alignment with Business Outcomes," discusses the need for segmentation, identification of stakeholders, and the guiding principles behind an organization's segmentation strategy. It also explores what happens when principles, strategies, and stakeholders change.

Chapter 3, "Developing a Segmentation Strategy," talks about developing a segmentation strategy, including determining key business requirements for segmentation and understanding how segmentation impacts the requirements of the business. In addition, it describes segmentation domains and lifecycles, including how to navigate how the need for segmentation changes over time, and covers static and dynamic approaches.

Chapter 4, "Macrosegmentation," discusses tactical implementation of segmentation at the macro level. It covers brownfield environment concepts such as VLANs, VRF instances, and dynamic enforcement mechanisms such as dynamic allocation into VLANs and creation of firewall rules.

Chapter 5, "Microsegmentation," discusses the challenges of microsegmentation and requirements to create a successful microsegmentation strategy across multiple network focus areas—campus, data center, remote access, and cloud environments.

Chapter 6, "Building the Segmentation Fabric," covers practical implementation steps for building a robust fabric capable of carrying layered segmentation methodologies. It sets the stage for future chapters on specific enforcement mechanisms for campus and data center environments that can be applied based on use case.

Chapter 7, "Implementing Segmentation with Cisco Technology," covers concepts related to identity and enforcement within the campus, including both macrosegmentation and microsegmentation approaches, and provides configuration examples.

Chapter 8, "Segmenting Applications in the Data Center and Cloud," discusses segmentation in data centers and cloud environments.

Chapter 9, "Validating Policies, Monitoring Enforcement, and Responding to Deviations," discusses verification of a segmentation strategy that has been implemented. It covers how to know how healthy a segmentation strategy is, what adjustments can be made and at what scale, and what products are best suited for monitoring.

Chapter 10, "Segmentation Maturity Model and Scorecard," presents a model segmentation maturity scorecard that you can use or adapt to measure your organization's goals and achievements.

Chapter 11, "Reference Architecture," summarizes how the concepts presented throughout the book can be applied to a network architecture to provide a reference for future implementations.

Finally, Chapter 12, "The Future of Segmentation," explores developing technologies, such as AI and centralized and integrated architectures. It also looks at the need to become even more granular in enforcement mechanisms to include kernel-level segmentation.

Summary

This chapter dives straight into the need for segmentation and the challenges that an organization commonly faces as it forms segmentation strategy. The chapter refers to the Zero Trust model discussed in the book *Zero Trust Architecture*, which is a standard approach to applying enforcement mechanisms to endpoints, workloads, and entities on a network after analyzing an organization's policies, identification mechanisms, and interaction capabilities. This chapter also discusses challenges an organization is likely to encounter in its segmentation strategy, such as the need for a robust and integrated vision across teams due to different stakeholder viewpoints. In concludes with a look ahead at the rest of this book.

Reference

- Cindy Green-Ortiz, et al., *Zero Trust Architecture*, Cisco Press, 2023.

Index

A

AAA (authentication, authorization, and accounting), 432

access control lists (ACLs), logical segmentation, 92–94

access layer, Cisco SAFE (Secure Architecture for Everyone), 82

ACI (Application Centric Infrastructure), 266–268

- APIC (Application Policy Infrastructure Controller), 291–292
- application profile (AP), 293
- application-centric topology, 293–295
- bridge domains (BD), 292
- Catalyst SD-WAN mesh, 298–299
- CNI (Container Network Interface), integration, 303–305
- contracts, 293, 307
- CSW (Cisco Secure Workload), 309–312
- data center, segmenting, 290–315
- endpoint group (EPG), 292, 306–308
- endpoint security group (ESG), 293, 307–308
- end-to-end segmentation, packet life, 312–315
- fabric overlay, 295–296
- fabric, 291–293
- L3Out (Layer 3 outside), 296–297
- network-centric topology, 293–295
- SD-Access, 297–298
- security inspections, 305–306
- service graphs, 305–306
- spine switches, 291
- tenants, 292
- VMM (Virtual Machine Manager), integration, 299–303
- VXLAN (Virtual Extensible LANs), 295–296

ACLs (access control lists)

- DACLs (downloadable ACLs), 93
- logical segmentation, 92–94
- macrosegmentation, 124–125

actionable framework, segmentation, 66–68

adaptability, hybrid networks, 42–43, 58–59

adaptive access, contextual awareness, 440–442

agentless workloads, CSW (Cisco Secure Workload), 278–279
agents, 421–422
CSW (Cisco Secure Workload), 277–278, 279
AI (Artificial Intelligence)
CSW policy discovery, 287–289
endpoint analytics, 259–262
future role, 438–440
hybrid networks, 29–30
resource utilization, 42
security risks, 32
AI-driven behavior analysis control (OSI), 389
alarms, Cisco Secure Network Analytics, 330–334
alerts, Cisco Secure Network Analytics, 330–334
algorithms, graph, 167
Amazon Web Services (AWS), 105–106
AMDB (asset management database), 421
analytics
endpoints, 259–262
reference architecture, 433–434
Zero Trust, 10–11
anomaly alarm, Cisco Secure Network Analytics, 333
APIC (Application Policy Infrastructure Controller), 291–292
Application Centric Infrastructure (ACI). *see* ACI (Application Centric Infrastructure)
application firewalls control (OSI), 387
application profile (AP), ACI (Application Centric Infrastructure), 293
application segmentation, 153–154
application team roles, hybrid networks, 65–66
application-centric topology, ACI (Application Centric Infrastructure), 293–295
applications
awareness, microsegmentation, 176–177
hybrid networks, 64–65
development, 63–64
mapping, 65
three-tier web, 157
architectural layers
Cisco SAFE (Secure Architecture for Everyone)
access layer, 82
business use cases layer, 81
core layer, 82
distribution layer, 82
endpoints layer, 82
services layer, 82
architecture. *see also* reference architecture
Cisco Secure Network Analytics, 320–322
hybrid networks, 56
collaboration and communication, 59–60
design and planning, 56–58
innovation and adaptation, 58–59
reference, 395–399
branch domains, 399–410
campus domains, 399–410
security onion, 396
Zero Trust, 175–176
architecture phase, Cisco SAFE (Secure Architecture for Everyone), 77–82

Armis, 273
Artificial Intelligence (AI). *see* AI (Artificial Intelligence)
ASA access lists, 93
asset classification control (OSI), 386
asset inventory management, 61
OSI (Open System Interconnection) model, 363, 366
reference architecture, 431–432
asset management
hybrid networks, 60–61
OSI (Open System Interconnection) model, 365–366
asset management database (AMDB), 421
attack surface segmentation score (CSW), 290
authentication modes (VLANs), 88–90
authorization policies, configuring, SGTs (Security Group Tags), 256
automated segmentation management control (OSI), 389
automation, 450
hybrid networks, 64–65
leveraging, 66
security risks, 32
troubleshooting, 66
unpredictable speed, 42
automation platforms, 365
AVC (Application Visibility and Control), Firewall Management Center, 231–233
AWS (Amazon Web Services), 105–106
AWS VPC flow logs, 323
Azure NSG2 flow logs, 323

B

Baelund, Peter, 54
behavior analytics, 433
border nodes, SD-Access, 204–207
branch domains, reference architecture, 399–410
branch networks, payment processing, protecting, 137–139
branches, networks, 77
bridge domains (BD), ACI (Application Centric Infrastructure), 292
bridge networking, containers, 107
Bring Your Own Device (BYOD). *see* BYOD (Bring Your Own Device)
building management endpoints, 409
business entities, 18–19
business flows, Cisco SAFE (Secure Architecture for Everyone), 74, 75–76
business unit-focused endpoints, 409
business use cases layer, Cisco SAFE (Secure Architecture for Everyone), 81
BYOD (Bring Your Own Device)
bandwidth issues, 42
hybrid networks, 27–28
security risks, 32

C

campus domains, reference architecture, 399–410
campuses
microsegmentation, 152–153
networks, 78
capability phase, Cisco SAFE (Secure Architecture for Everyone), 75–76

CAs (certificate authorities), 432
Catalyst Center, 190–191
Catalyst SD-WAN mesh, ACI (Application Centric Infrastructure), 298–299
chief information security officer (CISO), 358
CISA (Cybersecurity Information Sharing Act), 33
Cisco Catalyst Center, 190–191
Cisco Cyber Vision, 270–271
Cisco Flow Sensor, 320
Cisco Identity Services Engine, 192–194
integration, 186–187
Cisco Platform Exchange Grid (pxGrid), 335
Cisco SAFE (Secure Architecture for Everyone), 71–75
architectural layers
access layer, 82
business use cases layer, 81
core layer, 82
distribution layer, 82
endpoints layer, 82
services layer, 82
architecture phase, 77–82
business flows, 74, 75–76
capability phase, 75–76
design phase, 82–84
domains, 79
identifying segmentation needs, 424–428
Key to SAFE, 74, 78–79
macrosegmentation, mapping, 136
PINs, 77, 80
threat catalog, 74
toolkit, 74–75
Cisco Secure Firewall (CSF). *see* CSF (Cisco Secure Firewall)
Firewall Management Center
AVC (Application Visibility and Control), 231–233
Cisco Secure Workload, 237
configuring downloadable ACLs (DACLs), 246–255
configuring Dynamic VLAN, 245
configuring policies, 255–256
configuring switches, 244–245
context-aware policies, 237–240
endpoint analytics, 259–262
endpoint profiling, 256–258
ISE (Identity Services Engine), 235–237, 240–244
user-identity based policies, 233–235
Cisco Secure Network Analytics, 318–319
adoption phases, 339–340
alarms, 330–334
alerts, 330–334
architecture, 320–322
CSEs (custom security events), 334–335
CSW (Cisco Secure Workload), 340–344
flow ingestion, 322–327
group policy monitoring, 335–336
host groups, 327–328
observations, 330–334
Response Management, 336–339
security events, 330–334
SMC (Secure Management Console), 320
telemetry, 319
TrustSec policy analytics report, 336

Cisco Secure Workload, 421
Firewall Management Center, 237
Cisco Secure Workload (CSW). ***see*** **CSW (Cisco Secure Workload)**
Cisco Splunk, integration, 187
Cisco Telemetry Broker, 321, 326
Cisco Validated Designs (CVDs), 73
CISO (chief information security officer), 358
cloud, logical segmentation, as service in public, 102–104
cloud adoption, hybrid networks, 44–45
cloud instances, hybrid networks, 44
cloud locations, networks, 78
cloud PINs, workload communication security, 142–143
cloud-based security controls, 224–229
cloud-native segmentation, 106–111, 451
controls, 154–171
CMBDs (configuration management databases), 269–270
CMD (Cisco Meta Data), 95–96
CMDB (configuration management database), 421
CMM (Capability Maturity Model), SSMM (Secure Segmentation Maturity Model), 348–351
CNI (Container Network Interface), ACI (Application Centric Infrastructure), integration, 303–305
common services, scope hierarchy, 169
community VLANs, 91
compatibility, system, 58
compliance mandates
hybrid networks, 32–41
PCI DSS (Payment Card Industry Data Security Standard), 35–36
compliance segmentation score (CSW), 290
concern index alarm, Cisco Secure Network Analytics, 333
configuration management database (CMDB), 421
configuration management databases (CMDBs), 269–270
configuration segmentation score criteria, 383
containers, 107–108
dockers, 108
inter-pod communication, 109
intra-pod communication, 109
labels, 110
selectors, 110
context-aware policies
enforcement, 224–229
Firewall Management Center, 237–240
microsegmentation, 179
contextual awareness, adaptive access, 440–442
contracts, ACI (Application Centric Infrastructure), 293, 307
control plane (CP) node, SD-Access, 201–203
controls, hybrid networks, 46
core layer, Cisco SAFE (Secure Architecture for Everyone), 82
corporate endpoints, 408
CP (control plane) node, SD-Access, 201–203

create, read, update, and delete (CRUD) operations, 174
critical assets, identifying, 56
CRM (customer relationship management) software, 176
cross-team collaborations, 17–18
CRUD (create, read, update, and delete) operations, 174
CSEs (custom security events), Cisco Secure Network Analytics, 334–335
CSF (Cisco Secure Firewall), microsegmentation, 179–182
CSW (Cisco Secure Workload)
ACI (Application Centric Infrastructure), 309–312
agentless workloads, 278–279
agents, 277–278, 279
Cisco Secure Network Analytics, 340–344
filters, 282
first-time user experience, 276–280
measuring segmentation scores, 289–290
network traffic flow visibility, 280–283
pod vulnerability detection, 283
policies, 284–287
automated discovery with AI, 287–289
scope hierarchy, 281
Workload Traffic dashboard, 280
workloads
packages, 284–285
profiles, 283–285
Zero Trust, microsegmentation, 275–290
culture layer (network design), 4
custom security events (CSEs), Cisco Secure Network Analytics, 334–335
CVDs (Cisco Validated Designs), 73
Cybersecurity Information Sharing Act (CISA), 33

D

DACLs (downloadable ACLs), 93
configuring, 246–255
data centers
networks, 78
remote access security, 139–140
segmentation, ACI (Application Centric Infrastructure), 290–315
data classification control (OSI), 386
data hoarding alarm, Cisco Secure Network Analytics, 333
data plane (SD-Access), 195–196
data privacy, 444–445
Data Store, 321
deployment, 322
DDoS Source alarm, Cisco Secure Network Analytics, 333
DDoS Target alarm, Cisco Secure Network Analytics, 333
dependencies, entities, 8
design, hybrid networks, 58
design phase, Cisco SAFE (Secure Architecture for Everyone), 82–84
desktop, engineering, 15
DHCP (Dynamic Host Control Protocol), 134
digital experience monitoring, 128–129
Digital Operational Resilience Act (DORA), 33

discovery, policies, 174–175
distributed macrosegmentation, hybrid mesh firewalls, 445–449
distributed microsegmentation, hybrid mesh firewalls, 445–449
distribution layer, Cisco SAFE (Secure Architecture for Everyone), 82
DLPs (data loss prevention), 128–129
DNS security, reference architecture, 433
dockers, containers, 108
domains
branch, reference architecture, 399–410
campus, reference architecture, 399–410
Cisco SAFE (Secure Architecture for Everyone), 79
DORA (Digital Operational Resilience Act), 33
downloadable ACLs (DACLs), 93
configuring, 246–247
Dynamic Host Control Protocol (DHCP), 134
Dynamic VLAN, configuring, 245

E

EBPF (Extended Berkeley Packet Filter), 117
edge nodes (ENs), SD-Access, 203–204
edge PINs, IoT (Internet of Things) endpoints, 143–145
edge sprawl, hybrid networks, 44
edges, networks, 78
effort versus value matrix, hybrid networks, 48–50
Embedded Wireless Controller, SD-Access, 210
endpoint group (EPG), ACI (Application Centric Infrastructure), 292, 306–308
endpoint security group (ESG), ACI (Application Centric Infrastructure), 293, 307–308
endpoints, 13–14, 18–19
analytics, 259–262
building management, 409
business unit-focused, 409
corporate, 408
granular enforcement mechanisms, 182–186
macrosegmentation, 134–135
profiling, 256–258
research and development endpoints, 409
security, 409
endpoints layer, Cisco SAFE (Secure Architecture for Everyone), 82
end-to-end segmentation
ACI (Application Centric Infrastructure), packet life, 312–315
SD-Access, 219–224
end-user devices, 61
OSI (Open System Interconnection) model, 366
end-user management, OSI (Open System Interconnection) model, 363–364
enforcement, Zero Trust, 9–10
engineering
desktop, 15
network, 15
entities, 6
expected lifecycle, 7
identifying attributes, 7
known dependencies, 8

last evaluation date, 8
managers, 6
owners, 6
support team, 7
environment label, workloads, 156
ETA-capable devices, 323
events, Cisco Secure Network Analytics, 330–334
executive-driven vision, 19
exfiltration alarm, Cisco Secure Network Analytics, 333
expected lifecycles, entities, 7
exploitation alarm, Cisco Secure Network Analytics, 333
Extended Berkeley Packet Filter (EBPF), 117
extended nodes, SD-Access, 208–209

F

fabric, 189–190
ACI (Application Centric Infrastructure), 291–293
SD-Access, 190, 212–213
border nodes, 204–207
Cisco Catalyst Center, 190–191
Cisco Identity Services Engine, 192–194
control plane (CP) node, 201–203
edge nodes (ENs), 203–204
Embedded Wireless Controller, 210
end-to-end segmentation, 219–224
extended nodes, 208–209
Fabric in a Box, 210, 212
fabric roles, 201–202
fabric site, 212
fabric-mode access points, 210
intermediate nodes, 204
large site, 215–216
medium site, 214–215
operational planes, 195–196
overlay network, 198–200
reference model to deploy for distributed campus, 217–218
shared services, 200–201
small site, 213–214
transit control plane nodes, 211
transit networks, 211
underlay network, 197–198
WLCs (wireless LAN controllers), 209–210
Fabric in a Box, SD-Access, 210, 212
fabric site, SD-Access, 212
fabric-mode access points, SD-Access, 210
Federal Information Security Management Act (FISMA), 33
filters, CSW (Cisco Secure Workload), 282
Firepower Management Center (FMC), 179–180
Firewall Management Center, 231
AVC (Application Visibility and Control), 231–233
Cisco Secure Workload, 237
context-aware policies, 237–240
downloadable ACLs (DACLs), configuring, 246–255
Dynamic VLAN, configuring, 245
endpoints
analytics, 259–262
profiling, 256–258
ISE (Identity Services Engine), 235–237, 240–244
policies, configuring, 255–256

switches, configuring, 244–245
user-identity based policies, 233–235
firewall rules control (OSI), 387
firewalls
CSF (Cisco Secure Firewall), microsegmentation, 179–182
hybrid mesh, distributed segmentation, 445–449
IOS zone-based configuration, 98–99
logs, 323
macrosegmentation, 123–124
hybrid mesh, 127–130
interface-based, 125–126
zone-based, 126–127
next-generation, microsegmentation, 175–176
reference architecture, 433
FISMA (Federal Information Security Management Act), 33
five pillars of Zero Trust, 1–2
analytics, 10–11
enforcement, 9–10
identity, 5–8
policy and governance, 2–5
vulnerability management, 8–9
flow ingestion, Cisco Secure Network Analytics, 322–327
FMC (Firepower Management Center), 179–180
integration, 187–188
forensic segmentation score (CSW), 290

G

GDPR (General Data Protection Regulation), 39–41
General Data Protection Regulation (GDPR), 39–41
granular enforcement mechanisms, endpoints, 182–186
graph of communication, workloads, 167
Green-Ortiz, Cindy, 1–2
group policy monitoring, Cisco Secure Network Analytics, 335–336
grouping workloads, 162–164
communication patterns, 168–169
guest wi-fi security, macrosegmentation, 136–137

H

"hacker's playground," 30–32, 34
handshakes, QUIC (Quick UDP Internet Connections), 104
headers, VXLAN (Virtual Extensible LANs), 101
Health Insurance Portability and Accountability Act (HIPAA), 33, 37–38
hierarchical scope trees, 165
HIPAA (Health Insurance Portability and Accountability Act), 33, 37–38
host groups, Cisco Secure Network Analytics, 327–328
hybrid mesh firewalls, 127–130
distributed segmentation, 445–449
hybrid networks
adaptability, 42–43
administrative controls, 45
AI (Artificial Intelligence), 29–30
application team roles, 65–66
applications, 64–65
development, 63–64
architecture, 56
collaboration and communication, 59–60

design and planning, 56–58
innovation and adaptation, 58–59
integration and alignment, 58
asset management, 60–61
automation, 64–65
BYOD (Bring Your Own Device), 27–28
cloud adoption, 44–45
cloud instances, 44
compliance mandates, 32–41
GDPR (General Data Protection Regulation), 39–41
HIPAA (Health Insurance Portability and Accountability Act), 37–38
PCI DSS (Payment Card Industry Data Security Standard), 35–36
diverse tech stacks, 43
edge sprawl, 44
infrastructure, 64–65
IoT (Internet of Things), 28–29
leadership, 46–47
continuous improvement, 55–56
effort versus value matrix, 48–50
performance metrics, 54–55
resource allocation, 50–51
risk management, 52–53
stakeholder engagement, 53–54
strategic oversight, 47
training, 54
monitoring, 44
operational complexity, 43–44
organization, 46
patterns, 27
performance issues, 42–43
physical controls, 46
policies, 44
risk reduction, 46
scalability, 42–43
security, administration, 61–63
security risks, 30–32
segmentation, 25–26
actionable framework, 66–68
strategy development, 45–68
technical controls, 46

I

identifying attributes, entities, 7
identity, Zero Trust, 5–8
identity management segmentation score criteria, 383
Identity Services Engine, 192–194
Identity Services Engine (ISE), integration, 187–188
incident management, OSI (Open System Interconnection) model, 382–383
incident response, 64, 66
information security, 14–15
infrastructure
devices, 61
hybrid networks, 64–65
infrastructure device management, OSI (Open System Interconnection) model, 364–365
inspections, ACI (Application Centric Infrastructure), 305–306
instances, VRF (virtual routing and forwarding), 120–123

integrated Cisco solutions, 229–230
integration, 451
Cisco Cyber Vision, 270–271
microsegmentation, 186–188
OSI (Open System Interconnection) model, 370–372
security risks, 32
third-party, 271–273
interface-based firewalls, 125–126
intermediate nodes, SD-Access, 204
Internet of Things (IoT). *see* IoT (Internet of Things)
inter-pod communication, containers, 109
intra-pod communication, containers, 109
intra-VLAN protection domain, 406
intrusion prevention, 128–129
IOS zone-based firewall configuration, 98–99
IoT (Internet of Things)
data flooding, 42
endpoint security, edge PINs, 143–145
hybrid networks, 28–29
security risks, 32
IPAM (IP Address Management), 432
IPFIX, 323
ISE (Identity Services Engine), 323
ACI (Application Centric Infrastructure), 266–268
configuring switches for, 244–245
consistency, 263
Firewall Management Center, 235–237, 240–244
integration, 187–188
ISE common policy, multidomain segmentation, 263–265
isolated VLANs, 91
isolation, enhanced security, 64

K

kernel-level segmentation with eBPF, 112–113
key performance indicators (KPIs). *see* KPIs (key performance indicators)
Key to SAFE, 78–79
known dependencies, entities, 8
KPIs (key performance indicators), 10
identifying, 54–55
Kubernetes networking model, 110–111
architecture, 155
clusters, 170–171
containers, 157
policies, 160
workloads, 155

L

L3Out (Layer 3 outside), ACI (Application Centric Infrastructure), 296–297
labels
containers, 110
workloads, 156
large language models (LLMs), 29–30
large site, SD-Access, 215–216
last evaluation date, entities, 8
lateral attack vector, 30–31, 32
leadership
hybrid networks, 46–47
effort versus value matrix, 48–50

performance metrics, 54–56
resource allocation, 50–51
risk management, 52–53
stakeholder engagement, 53–54
strategic oversight, 47
training, 54
leaf switches, ACI (Application Centric Infrastructure), 291
least privilege, 442–443
lifecycle management, OSI (Open System Interconnection) model, 363, 366
lifecycles
entities, 7
management, 61
Linux, kernel programmability with eBPF, 112
LLMs (large language models), 29–30
logging, reference architecture, 433–434
logical segmentation, 87
ACLs (access control lists), 92–94
cloud-native, 106–111
kernel-level segmentation with eBPF, 112–113
network virtualization, 99–100
PVLANs (private virtual local-area networks), 90–92
QUIC (Quick UDP Internet Connections), 102–104
SD (software-defined), 94–99
as service in public cloud, 105–106
VLANs (virtual local-area networks), 87–89
VXLANs (Virtual Extensible LANs), 101–102
wireless SSIDs (service set identifiers), 92
logs, firewalls, 323

M

macrosegmentation, 5–6, 117
ACLs (access control lists), 124–125
distributed, 445–449
hybrid mesh firewalls, 445–449
endpoints, 134–135
firewalls, 123–124
hybrid mesh, 127–130
interface-based, 125–126
zone-based, 126–127
guest wi-fi security, 136–137
IoT (Internet of Things) endpoints, edge PINs, 143–145
mapping, 136
payment processing, branch networks, 137–139
policies, 130–134
remote access security, data centers, 139–140
SaaS application access, PINs, 140–141
security onion, 117–118
virtualization, 120–122
visibility, 118–119
VLANs (virtual local-area networks), 122–123
workload communication security, cloud PINs, 142–143
Zero, 1
management plane (SD-Access), 196
managers, entities, 6
mapping, macrosegmentation, 136
MASQUE (Multiplexed Application Substrate over QUIC Encryption), 103
maturity level reassessment, OSI (Open System Interconnection) model, 353

measuring, segmentation, 175
Medigate, 272
medium site, SD-Access, 214–215
microperimeters, nanosegmentation, 443–444
microsegmentation, 147–148
application, 153–154
application awareness, 176–177
campuses, 152–153
cloud-native controls, 154–171
context-aware policies, 179
CSF (Cisco Secure Firewall), 179–182
distributed, 445–449
hybrid mesh firewalls, 445–449
endpoints, granular enforcement mechanisms, 182–186
implementation, 148–151
integration, 186–188
measuring, 175
next-generation firewalls, 175–176
policies, 171–172
creating, 172–173
discovery, 174–175
SGTs (Security Group Tags), 178–179
user identity policies, 176–177
Zero Trust, CSW (Cisco Secure Workload), 275–290
Microsoft System Center Configuration Manager (SCCM), 11–12
mixed workloads, 42
MnT (Monitor and Troubleshooting) node, ISE (Identity Services Engine), 193
money layer (network design), 4
monitoring, hybrid networks, 44
monitoring segmentation score criteria, 383
multi-auth switch port authorization mode (VLANs), 89
multidomain segmentation, ISE common policy, 263–265
multi-domain switch port authorization mode (VLANs), 89
multi-host switch port authorization mode (VLANs), 89

N

namespaces, containers, 107
nanosegmentation, 450
microsegmentation, 443–444
NCCM (network configuration and change management) tools, 364–365
NetFlow, 322
network access control lists control (OSI), 387
Network and Information Security Directive (NIS2), 33
network anomaly segmentation score (CSW), 290
network configuration and change management (NCCM) tools, 364–365
network engineering, 15
network traffic flow visibility, CSW (Cisco Secure Workload), 280–283
network virtualization, logical segmentation, 99–100
Network Visibility Module (NVM) flow data, 323
network-centric topology, ACI (Application Centric Infrastructure), 293–295
networks. *see also* hybrid networks
branches, 77
campuses, 78

cloud locations, 78
data centers, 78
design layers, 3–4
edges, 78
security, 15–16
WANs (wide-area networks), 78
next-generation firewalls, microsegmentation, 175–176
NIS2 (Network and Information Security Directive), 33
NSEL (NetFlow Security Event Logging), 323

O

observations, Cisco Secure Network Analytics, 330–334
onboarding process, reference architecture, 431–432
on-premises segmentation, 451
operational planes, SD-Access, 195–196
operations, 16–17
orchestration platforms, 365
Ordr, 272
organization, workloads, 155–156, 164–165
OSI (Open System Interconnection) model, 351–352. ***see also*** **SSMM (Secure Segmentation Maturity Model)**
asset inventory management, 363, 366
asset management, 365–366
collaboration and communication, 374–375
configuration, 380
continuous improvement, 357
design and implementation, 379
design and planning, 368–370
end-user devices, 366
end-user management, 363–364
feature controls, 383–389
identity and access management, 382
incident management, 382–383
infrastructure device management, 364–365, 366
innovation and adaptation, 372–374
integration across hybrid environments, 380–381
integration and alignment, 370–372
lifecycle management, 363, 366
maintenance, 380
maturity level reassessment, 353
monitoring, 380
network administrator responsibilities, 379
optimization, 380
performance metrics, 357
policy development, 355
resource allocation, 354–355
risk management, 355
role-based, 353–354
scorecard, 358–360
architecture effectiveness, 375–377
network and security administrative effectiveness, 383
security administrator responsibilities, 381
security policy enforcement, 381
SSMM (Secure Segmentation Maturity Model), 360–362
stakeholder engagement, 356
strategic oversight, 354

threat assessments, 352–353
threat detection and response, 381–382
training, 356
overlay control plane (SD-Access), 195
overlay network, SD-Access, 198–200
overlay networks, containers, 107
owners, entities, 6

P

packages, workloads, 284–285
packets, end-to-end segmentation, 312–315
PAN (Policy Administration Node), ISE (Identity Services Engine), 193
Payment Card Industry Data Security Standard (PCI DSS), 33, 35–36
payment processing, branch networks, protecting, 137–139
PCI DSS (Payment Card Industry Data Security Standard), 33, 35–36
PDPs (policy decision points), 134–135
performance issues, hybrid networks, 42–43
performance metrics
hybrid networks, 54–55
OSI (Open System Interconnection) model, 357
physical segmentation, 86
pillars of Zero Trust, 1–2
analytics, 10–11
enforcement, 9–10
identity, 5–8
policy and governance, 2–5
vulnerability management, 8–9
PINs
Cisco SAFE (Secure Architecture for Everyone), 77, 80
cloud, workload communication security, 142–143
SaaS application access, 140–141
PIPs (policy information points), 134–135
pod vulnerability detection, CSW (Cisco Secure Workload), 283
policies
authorization, SGTs (Security Group Tags), 256
context-aware, 179, 237–240
enforcement, 224–229
CSW (Cisco Secure Workload), 284–287
automated discovery with AI, 287–289
group policy monitoring, Cisco Secure Network Analytics, 335–336
hybrid networks, 44
information security, 14–15
Kubernetes networking model, 160
macrosegmentation, 130–134
microsegmentation, 171–172
creating, 172–173
discovery, 174–175
PDPs (policy decision points), 134–135
PIPs (policy information points), 134–135
updates, 66

user identity, 176–177
user-identity based, 233–235
policy and governance, Zero Trust, 2–5
policy decision points (PDPs), 134–135
policy development, OSI (Open System Interconnection) model, 355
policy information points (PIPs), 134–135
policy plane (SD-Access), 196
policy violation alarm, Cisco Secure Network Analytics, 333
policy-based network management, 317–318
politics layer (network design), 3
port mapping, containers, 107
process hash segmentation score (CSW), 289
profiles, workloads, 283–285
profiling endpoints, 256–258
propagation process, SGTs (Security Group Tags), 265–266
provisioning segmentation score criteria, 383
proxies, reference architecture, 433
proxy logs, 323
PSN (Policy Service Node), ISE (Identity Services Engine), 194
public cloud integrations, 268–269
Purdue Model to segment industrial networks, 28
PVLANs (private virtual local-area networks), logical segmentation, 90–92
pxGrid (Cisco Platform Exchange Grid), 335
pxGrid (Platform Exchange Grid) node, ISE (Identity Services Engine), 194

Q

QUIC (Quick UDP Internet Connections)
handshakes, 104
logical segmentation, 102–104
MASQUE (Multiplexed Application Substrate over QUIC Encryption), 103

R

Raspberry Pi devices, lateral attack vector, 30–31
recon alarm, Cisco Secure Network Analytics, 333
reference architecture, 395–399, 417–418
AAA (authentication, authorization, and accounting), 432
analytics, 433–434
asset inventory management, 431–432
asset monitoring and discovery, 431–432
behavior analytics, 433
beyond branch, 419–424
beyond campus, 419–424
beyond endpoint focus, 415–416
branch domains, 399–410
campus domains, 399–410
CAs (certificate authorities), 432
configuration management database (CMDB), 431–432
DNS security, 433
example, 411–415

firewalls, 433
IPAM (IP Address Management), 432
logging, 433–434
onboarding process, 431–432
previous architectures, 430–431
proxies, 433
traffic visibility, 433
reference model to deploy for distributed campus, SD-Access, 217–218
region label, workloads, 156
remote access security, data centers, 139–140
research and development endpoints, 409
resource allocation
hybrid networks, 50–51
OSI (Open System Interconnection) model, 354–355
Response Management, Cisco Secure Network Analytics, 336–339
risk management
hybrid networks, 52–53
OSI (Open System Interconnection) model, 355
role-based perspectives, OSI (Open System Interconnection) model, 353–354
roles (fabric), SD-Access, 201–202
routers, multiple VRF tables, 100

S

SaaS, application access, PINs, 140–141
SAFE (Secure Architecture for Everyone). ***see*** **Cisco SAFE (Secure Architecture for Everyone)**
Sampled NetFlow, 323
scalability, hybrid networks, 42–43
SCCM (System Center Configuration Manager), 11–12
scope discovery, automating, 166–169
scope hierarchy, CSW (Cisco Secure Workload), 281
scorecard
OSI (Open System Interconnection) model, 358–360
architecture effectiveness, 375–377
network and security administrative effectiveness, 383
SSMM (Secure Segmentation Maturity Model), 360–362
evaluation criteria, 367
mapping, 377–378
SD (software-defined) segmentation, 94–97
SDA sites, increasing size, 400
SD-Access, 190
ACI (Application Centric Infrastructure), 297–298
border nodes, 204–207
Cisco Catalyst Center, 190–191
Cisco Identity Services Engine, 192–194
control plane (CP) node, 201–203
design, 212–213
edge nodes (ENs), 203–204
Embedded Wireless Controller, 210
end-to-end segmentation, 219–224
extended nodes, 208–209
Fabric in a Box, 210, 212
fabric roles, 201–202
fabric site, 212
fabric-mode access points, 210
intermediate nodes, 204
large site, 215–216
medium site, 214–215

operational planes, 195–196
overlay network, 198–200
reference model to deploy for distributed campus, 217–218
shared services, 200–201
small site, 213–214
transit control plane nodes, 211
transit networks, 211
underlay network, 197–198
WLCs (wireless LAN controllers), 209–210
Secure Architecture for Everyone (SAFE). *see* Cisco SAFE (Secure Architecture for Everyone)
Secure Network Analytics, integration, 186–187
Secure Segmentation Maturity Model (SSMM). *see* SSMM (Secure Segmentation Maturity Model)
Secure Workload, integration, 188
security
attack surface, 64
Cisco Secure Network Analytics, 318–319
architecture, 320–322
telemetry, 319
data privacy, 444–445
hybrid networks
administration, 61–63
risks, 30–32
incident response, 64, 66
isolation, 64
lateral attack vector, 30–31
security endpoints, 409
security events, Cisco Secure Network Analytics, 330–334
Security Group Tag Exchange Protocol (SXP), 265–266
Security Group Tags (SGTs). *see* SGTs (Security Group Tags)
security information and event management (SIEM) systems, 365
security onion, 117–118, 396
security policies and governance control (OSI), 389
security zones, logical segmentation, 97–99
security zones control (OSI), 387
segmentation, 1–2, 85. *see also* hybrid networks
AI (Artificial Intelligence), future role, 438–440
automation, 450
boundaries, defining, 56
core elements, 437
end-to-end, 219–224
fabric, 189–190
SD-Access, 190
failed strategies, 11–21
hybrid networks, 25–26
actionable framework, 66–68
adaptability, 42–43
cloud adoption, 44–45
compliance mandates, 32–41
leadership, 46–56
operational complexity, 43–44
performance issues, 42–43
scalability, 42–43
strategy development, 45–68
logical, 87
ACLs (access control lists), 92–94
cloud-native, 106–111
kernel-level segmentation with eBPF, 112–113
network virtualization, 99–100

PVLANs (private virtual local-area networks), 90–92
QUIC (Quick UDP Internet Connections), 102–104
SD (software-defined), 94–97
security zones, 97–99
as service in public cloud, 105–106
VLANs (virtual local-area networks), 87–89
VXLANs (Virtual Extensible LANs), 101–102
wireless SSIDs (service set identifiers), 92
macrosegmentation, 1, 117
ACLs (access control lists), 124–125
distributed, 445–449
endpoints, 134–135
firewalls, 123–124, 125–130
guest wi-fi security, 136–137
IoT (Internet of Things) endpoints, 143–145
mapping, 136
payment processing, 137–139
policies, 130–134
remote access security, 139–140
SaaS application access, 140–141
security onion, 117–118
virtualization, 120–122
visibility, 118–119
VLANs (virtual local-area networks), 122–123
workload communication security, 142–143
microsegmentation, 147–148
application, 153–154
application awareness, 176–177
automated zero trust, 160–171
campuses, 152–153
cloud-native controls, 154–160
context-aware policies, 179
CSF (Cisco Secure Firewall), 179–182
distributed, 445–449
granular enforcement mechanisms, 182–186
implementation, 148–151
integration, 186–188
measuring, 175
next-generation firewalls, 175–176
policies, 171–173, 174–175
SGTs (Security Group Tags), 178–179
user identity policies, 176–177
multidomain, ISE common policy, 263–265
nanosegmentation, 443–444, 450
overuse, 11
physical, 86
SD-Access, 190
shared responsibility model, 113–115
stakeholders, 3
strategy development, 71–72
Cisco SAFE (Secure Architecture for Everyone), 72–84
segmentation scores, measuring, CSW (Cisco Secure Workload), 289–290
selectors, containers, 110
service graphs, ACI (Application Centric Infrastructure), 305–306

services layer, Cisco SAFE (Secure Architecture for Everyone), 82
sFlow, 323
SGTs (Security Group Tags), 95
authorization policies, 256
microsegmentation, 178–179
propagation process, 265–266
shared responsibility model, segmentation, 113–115
shared services, SD-Access, 200–201
SIEM (security information and event management) systems, 365
single-host switch port authorization mode (VLANs), 88–89
small site, SD-Access, 213–214
SMC (Secure Network Analytics Management Console), 320
software-defined (SD) segmentation, 94–97
software-defined networking control (OSI), 388
spine switches
ACI (Application Centric Infrastructure), 291
leaf switches, 291
SSE (security service edge), 224–229
SSIDs (service set identifiers), logical segmentation, 92
SSMM (Secure Segmentation Maturity Model), 347–348. ***see also*** **OSI (Open System Interconnection) model**
CMM (Capability Maturity Model), 348–351
OSI (Open System Interconnection) model, 360–362
scorecard
evaluation criteria, 367
mapping, 377–378
mapping evaluation criteria, 390–392
stakeholder engagement, OSI (Open System Interconnection) model, 356
stakeholders
educating, 66
hybrid networks, engagement, 53–54
segmentation, 3
strategic oversight, OSI (Open System Interconnection) model, 354
strategies
aligning, 21–22
failed, 11–21
maturation, 22–23
support teams, entities, 7
switches, configuring, 244–245
switch-port authorization modes (VLANs), 88–90
SXP (Security Group Tag Exchange Protocol), 265–266
Syslog, 323

T

target index alarm, Cisco Secure Network Analytics, 333
tech stacks, hybrid networks, 43–44
telemetry
Cisco Secure Network Analytics, 319
protocols, 322–324
templates, policy, CSW (Cisco Secure Workload), 287
tenants
ACI (Application Centric Infrastructure), 292
VRF (virtual routing and forwarding) instances, 292
third-party integrations, 271–273

threat assessments, OSI (Open System Interconnection) model, 352–353
threat catalog, Cisco SAFE (Secure Architecture for Everyone), 74
threat detection and response, OSI (Open System Interconnection) model, 381–382
three-tier web application, 157
toolkit, Cisco SAFE (Secure Architecture for Everyone), 74–75
traffic analysis domain, 405
traffic visibility, reference architecture, 433
training
hybrid networks, 54
OSI (Open System Interconnection) model, 356
transit control plane nodes, SD-Access, 211
transit networks, SD-Access, 211
TrustSec, 95
classification process, 96
configuring, 247–255
enforcement process, 97
policy analytics report, Cisco Secure Network Analytics, 336
propagation process, 96–97

U

underlay network, SD-Access, 197–198
unified policy enforcement, workload context, 268–269
use cases, Cisco SAFE (Secure Architecture for Everyone), 75–76
user identity policies, microsegmentation, 176–177
user-identity based policies, Firewall Management Center, 233–235

V

veth pairs, containers, 107
Virtual Extensible LANs (VXLANs), 101–102
Virtual Machine Manager (VMM). ***see*** **VMM (Virtual Machine Manager)**
virtual machines, 107
virtual network interfaces, containers, 107
virtual private clouds control (OSI), 388
virtual routing and forwarding (VRF). ***see*** **VRF (virtual routing and forwarding)**
virtualization
logical segmentation, 99–100
macrosegmentation, 120–122
visibility, macrosegmentation, 118–119
VLAN control (OSI), 386
VLAN protection domains, 403
VLANs (virtual local-area networks)
authentication modes, 88–90
community, 91
isolated, 91
logical segmentation, 87–89
macrosegmentation, 122–123
VMM (Virtual Machine Manager), ACI (Application Centric Infrastructure), integration, 299–303
VPC label, workloads, 156

VRF (virtual routing and forwarding), 11, 118
instance protection domain, 404
instances, 120–123
ACI (Application Centric Infrastructure), 292
vulnerability management, Zero Trust, 8–9
vulnerability management and compliance scanning tools, 365
vulnerability segmentation score (CSW), 289
VXLAN (Virtual Extensible LANs)
headers, 101
VTEPs (VXLAN tunnel endpoints), 101–102
VXLAN control (OSI), 386
VXLANs (Virtual Extensible LANs), logical segmentation, 101–102

W

WANs (wide-area networks), 78
wireless LAN controllers (WLCs), 131
wireless SSIDs (service set identifiers), logical segmentation, 92
WLCs (wireless LAN controllers), 131
SD-Access, 209–210
workflows, workload classification, 269
Workload Traffic dashboard (CSW), 280
workloads, 154
communication, securing, 142–143
graph of communication, 167
grouping, 162–164
communication patterns, 168–169
information, 170–171
Kubernetes networking model, 155
labels, 156
mapping context for unified policy enforcement, 268–269
organization, 155–156, 164–165
packages, 284–285
profiles, 283–285

Z

Zeek logs, 324
Zero Trust, 1–2, 11, 275
architecture, 175–176
CSW (Cisco Secure Workload), microsegmentation, 275–290
five pillars, 1–2
analytics, 10–11
enforcement, 9–10
identity, 5–8
policy and governance, 2–5
vulnerability management, 8–9
granular segmentation, 440–444
least privilege, 442–443
Zero Trust Network Access (ZTNA), 315
Zero Trust access control (OSI), 388
Zero Trust Architecture, **1–2**
zone-based firewalls, 126–127
ZTNA (Zero Trust Network Access), 315