

Designing Real-World Multi-Domain Networks



ciscopress.com

**DHRUMIL PRAJAPATI
JEREMY BOWMAN
NAVIN SUVARNA**

FREE SAMPLE CHAPTER |



Designing Real-World Multi-Domain Networks

Dhrumil Prajapati CCIE No. 28071, CCDE No. 20210002

Jeremy Bowman CCIE No. 51241, CCDE No. 20180016

Navin Suvarna CCIE No. 24583

Cisco Press

Hoboken, New Jersey

Designing Real-World Multi-Domain Networks

Dhrumil Prajapati, Jeremy Bowman, Navin Suvarna

Copyright© 2024 Cisco Systems, Inc.

Published by:
Cisco Press

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

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

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

\$PrintCode

Library of Congress Control Number: 2024903903

ISBN-13: 978-0-13-803721-5

ISBN-10: 0-13-803721-3

Warning and Disclaimer

This book is designed to provide information about how multiple domains inter-work with each other. The methods and architectural designs explained here are deployed in real networks and function effectively. Every effort has been made to make this book as complete and as accurate as possible, but no warranty or fitness is implied. Just like with any large network deployments, the authors of this book highly recommend testing the integrations and its functionality in a lab environment and running all your organization-specific test cases before production deployment.

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 authors 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.

Special Sales

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

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

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

Feedback Information

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

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

We greatly appreciate your assistance.

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

GM K12, Early Career and Professional Learning: Soo Kang

Alliances Manager, Cisco Press: Caroline Antonio

Director, ITP Product Management: Brett Bartow

Managing Editor: Sandra Schroeder

Development Editor: Ellie C. Bru

Senior Project Editor: Mandie Frank

Copy Editor: Chuck Hutchinson

Technical Editors: Arul Jagadeesan, Kyle Barnes

Editorial Assistant: Cindy Teeters

Designer: Chuti Prasertsith

Composition: codeMantra

Indexer: Ken Johnson

Proofreader: Jennifer Hinchliffe



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

Dhrumil Prajapati: Dhrumil is a principal architect within Cisco CX's GES Architecture team. His focus lies on multi-domain networks, and he has been offering a complete life-cycle of professional services and architecture advisory for the past 13 years. His expertise extends to serving enterprise, government, and service provider entities across the globe. His services are designed to assist clients in planning, designing, deploying, managing, and interoperating all networking technology domains within their private or public infrastructure and application environments.

In his networking career, Dhrumil has designed networks for more than 150 organizations, which inspired him to write a book on the subject. He is a coauthor of *Cisco SD-Access for Industry Verticals* (<https://cs.co/sda-verticals-book>), and holds patents and has given multiple presentations in Cisco Live on SD-Access and multi-domain.

Dhrumil holds dual CCIEs in Enterprise Infrastructure and Service Provider, as well as a CCDE, in addition to other leading technical certificates. He also assists the Cisco Certifications team by reviewing and providing feedback for Cisco certificate exams. In addition, he leads several initiatives within Cisco CX aimed at driving delivery standardization and enhancing efficiency through automation innovation.

Currently residing in Apex, North Carolina, Dhrumil has a passion for motor racing, woodworking, and innovative electronics that enhance human life. His wife, Devanshi, and son, Ram, bring pure joy to his life, adding a touch of fun to every day.

Jeremy Bowman: Jeremy is a senior solutions architect within Cisco CX's GES Architecture team. Within the GES team, he has more than 10 years of experience focusing on designing and implementing large-scale multi-domain IBN environments that meet the client-specific requirements. Additionally, he has presented several sessions at Cisco Live. Jeremy is a CCDE and double CCIE in Enterprise Infrastructure and Security. He also has professional-level certification in three other Cisco certification tracks. Inside Cisco, Jeremy is a mentor to other architects and engineers, sits on several CX design review panels, and interacts with engineering to help further product development. When not working, he enjoys spending time with his family at home, camping, or on the lake.

Navin Suvarna: Navin possesses more than 18 years of extensive experience as a seasoned network delivery architect. With a decade-long tenure at Cisco, he has adeptly architected, designed, tested, operationalized, managed, and optimized expansive networks across service providers, global enterprises, and public sector segments. Holding the distinguished CCIE (No. 24583) in enterprise networking, Navin boasts an impressive collection of professional-level certifications.

Since the inception of software-defined networking, Navin has actively contributed to Cisco's incubation team, playing a pivotal role in incubating multiple intent-based solutions. His expertise has been instrumental in facilitating the adoption, assimilation, and rapid deployment of intricate multi-domain solutions. In the past four years, Navin has held the role of lead delivery architect for the United States Public Sector team. While specializing in multi-domain solutions, his proficiency in aligning network infrastructure

with the evolving needs of the public sector has garnered affiliations with thought leaders in both the public sector and Cisco engineering. This collaboration has resulted in the development of novel features, functions, and validated practices tailored to multi-domain architectures, addressing the specific requirements of public sector clients.

While Navin's specialization remains rooted in multi-domain intent-based networks, his recent focus centers on educating clients about the value proposition of automation and guiding them on their automation journey. An active speaker, Navin has shared his insights and experiences of architecting multi-domain solutions on global platforms, including Cisco Live. He also champions hands-on demonstrations and automation show-cases, offering real-world insight into intent-based networking through automation use cases.

Beyond his professional pursuits, Navin finds joy in his 14-year marriage to his wife, Rupali. Their family includes their 12-year-old daughter, Trisha, and 6-year-old son, Vivaan. They currently reside in Research Triangle Park, North Carolina.

About the Technical Reviewers

Arul Jagadeesan (CCIE No. 9942) is a network architect in Cisco's Customer Experience Organization. As a network architect, Arul collaborates with customer IT organizations to provide a suite of professional consultative architectural and engineering solutions to achieve a high-performance network for the next generation of technology and business applications across a global infrastructure. Arul has presented at industry events such as Cisco Live on Software-Defined Networking, plus Security and Enterprise Architecture. He is PMP, CISSP, AWS, and Azure certified. Arul holds a bachelor of science degree in electronics and communication engineering and a master of science in telecommunications.

Kyle Barnes (CCIE No. 46535, CCNP DC, CCDP) is a customer delivery architect with Cisco and has more than 13 years of experience in the IT industry. He has worked in both the private and public sector alike, supporting different capacities ranging from network engineer to consultant to architect. It is with his real-world experience and aptitude for simplifying the complex that Kyle has developed a skillset of solving difficult networking scenarios. Kyle's philosophy of understanding the business requirements first has been his North Star when approaching any complex scenario. Most importantly, Kyle's humility (as he writes this in the third person) has made him an incredible relationship builder with both peers and customers alike.

Dedications

Dhrumil Prajapati:

Writing a book is a monumental task that requires a tremendous amount of personal and emotional support to reach fruition. The completion of this book would not have been possible without the unwavering support and boundless love of my beautiful wife, Devanshi. She has been my beacon, my guiding light, and the rockstar of my life. Her patience, understanding, and belief in my abilities have been the cornerstone of this journey.

My son, Ram, has shown me the true meaning of curiosity, inspiring me to explore new ideas rather than simply accepting them. He has been my source of motivation and a powerhouse of energy, propelling me to consistently strive for better.

I am deeply grateful for their sacrifices as I worked into the late hours of the night and over the weekends to bring this book to life. It's now time for us to celebrate this achievement together. This book stands as my first major accomplishment, and it was made possible because of them.

Jeremy Bowman:

Pauly, you left your family behind and came to a foreign country. You did not know what the future would hold, but you came willingly anyway. None of this would have been possible had you not been there every step of the way. There were many late-night and all-night maintenance windows, client visits, and unscheduled calls. There were lost holidays and other sacrifices so the work could get done. All of those events helped lead to the culmination of this work. You will never know how much you mean to me nor how much I appreciate all that you have done. It will always be demasiado.

Matthew, Megan, and Michael, you are my superheroes. Each of you inspire me to push harder to be a better version of me and to make you proud. I am so proud to be your dad.

Navin Suvarna:

To my esteemed parents, whose resolute support and endless encouragement have laid the foundation of my journey in the realm of technology. Your guidance and belief in my abilities have propelled me forward.

To my cherished wife, Rupali, your unwavering support, understanding, and encouragement have been my steadfast companions throughout this technical expedition. Your love has illuminated every intricate path I've traversed.

To my invaluable children, Trisha and Vivaan, your boundless energy and curiosity have inspired me to push the boundaries of knowledge. Your presence has imbued every technical endeavor with purpose and meaning.

This book is dedicated to you, my pillars of strength and inspiration, for your unrelenting support and for being the driving force behind me and for that I am forever indebted to you.

Acknowledgments

Dhrumil: I extend my heartfelt thanks to my colleagues, team, mentors, friends, and leaders. Their inspiration and collaboration have been invaluable in the creation of this book. This book is a testament to our shared journey, and I hope it will inspire others as much as they all have inspired me.

Jason Gooley, thank you for guiding me to take the first step for this book.

Jeremy: To my team, colleagues, leaders, and friends, you have all helped and contributed to this work. Thank you for your support and encouragement.

Olmedo and Abraham, you taught me more than you will ever know.

Navin: I extend my sincere gratitude to Arvind Chari and John Weston for their invaluable assistance with my ACI queries, and to Kevin Manweiler for helping me with my Meraki questions. Their expertise and insights greatly contributed to the quality of these chapters.

Additionally, I want to express my appreciation to my colleagues and friends at Cisco who patiently responded to my numerous inquiries throughout the completion of my chapters.

Reader Services

Register your copy at www.ciscopress.com/title/ISBN 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 9780138037215 and click Submit.

*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 at a Glance

Introduction xxiv

Part I Introduction

Chapter 1 Multi-Domain Networks 1

Part II Multi-Domain Design

Chapter 2 SD-Access and Campus Fabric 21

Chapter 3 SD-WAN and DMVPN 45

Chapter 4 Application Centric Infrastructure (ACI)—Integration and Multi-Domain Capabilities 67

Chapter 5 Enterprise MPLS 163

Chapter 6 Carrier Neutral Facilities 191

Chapter 7 Cloud 229

Chapter 8 Security 253

Chapter 9 Automation 273

Part III Real-World Use Cases

Chapter 10 Manufacturing Use Case with SDA, SD-WAN, and CNF 285

Chapter 11 Financial Use Case 307

Chapter 12 Retail Use Case Using CNF, SD-WAN, and ACI 327

Chapter 13 Public Sector Use Case 381

Chapter 14 Transportation Use Case 419

Index 441

Contents

	Introduction	xxiv
Part I	Introduction	
Chapter 1	Multi-Domain Networks	1
	Overview	1
	What Are Multi-Domain Networks?	2
	Components of Multi-Domain Networks	3
	<i>Redundancy</i>	3
	<i>Resiliency</i>	9
	<i>WAN Redundancy and Multi-Connectivity</i>	9
	<i>Shared Services</i>	10
	<i>Multiple Technologies</i>	11
	Characteristics of Multi-Domain Networks	11
	<i>Entry and Exit Points</i>	12
	<i>Trust Boundaries</i>	13
	<i>Different Areas</i>	13
	<i>Fault Domains</i>	13
	Why Do We Need Multi-Domain Networks?	14
	Information Technology as a Business (ITaaS)	14
	<i>Business Units</i>	15
	<i>IT Boundary</i>	16
	Multi-Tenancy	17
	<i>Traffic Isolation</i>	17
	<i>Network Segmentation</i>	18
	Automation and Orchestration	19
	Summary	19
Part II	Multi-Domain Design	
Chapter 2	SD-Access and Campus Fabric	21
	Overview	21
	Interaction with the Outside World	22
	Cisco Software Defined Access (SDA)	23
	Campus Fabric	25
	Inter-working with SD-WAN	25
	Default End-to-End Macrosegmentation	26

	Integrated and Nonintegrated SD-WAN Solutions	27
	Learnings from SD-Access and SD-WAN Integration	30
	Inter-working with ACI and Data Center Fabric	31
	Inter-working with MPLS	34
	Inter-working with CNF	37
	End-to-End Segmentation Using SDA-Transit	38
	Multi-Site Remote Border (MSRB)	40
	Inter-working with Security Stack	41
	Network Access Control (NAC)	42
	Perimeter Access	43
	Summary	44
Chapter 3	SD-WAN and DMVPN	45
	Overview	45
	SD-WAN	45
	SD-WAN and SDA	46
	One-Box SDA and SD-WAN	46
	Catalyst SD-WAN Manager and Catalyst Center Integration	48
	Two-Box SDA and SD-WAN	48
	SDA and SD-WAN Segmentation	49
	SDA and SD-WAN Best Practices	51
	SD-WAN and ACI	52
	Catalyst SD-WAN Manager and APIC Integration	55
	ACI and SD-WAN Segmentation	56
	ACI and SD-WAN Best Practices	56
	SD-WAN with MPLS	56
	SD-WAN and the Cloud	57
	SIG	58
	Cloud OnRamp	59
	DMVPN	61
	DMVPN and SDA	62
	DMVPN and SDA Policy Enforcement	64
	SDA and DMVPN Best Practices	64
	DMVPN and ACI	64
	ACI and DMVPN Segmentation	65
	ACI and DMVPN Best Practices	65
	Summary	66

Chapter 4	Application Centric Infrastructure (ACI)—Integration and Multi-Domain Capabilities	67
	Overview	67
	Key Components	70
	OpFlex Control Protocol	70
	ACI Physical Design	72
	<i>Application Policy Infrastructure Controller</i>	74
	<i>Spines and Leaf Fabric Design</i>	74
	<i>Nexus Dashboard Orchestrator</i>	75
	ACI Policy Model	77
	<i>Tenant Policies</i>	77
	<i>Access Policies</i>	79
	<i>Layer 4–7 Services</i>	80
	Routing and Forwarding in ACI Fabric	81
	VXLAN Overlay	81
	<i>Endpoint Learning and Council of Oracles Protocol (COOP)</i>	83
	ACI and SD-Access Pairwise Integration	88
	SD-Access–ACI: Group/Identity Automated Mapping Federation (Microsegmentation Use Case)	89
	SD-Access–ACI: VN-to-VRF Automated Mapping Federation (Macrosegmentation Use Case)	93
	SD-Access–ACI: VN-to-VRF Manual Mapping (Macrosegmentation Use Case)	96
	Guidelines and Limitations Around ACI–SD-Access Integration	101
	ACI and SD-WAN Pairwise Integration	101
	<i>DC and Branch Using SD-WAN</i>	104
	<i>Multi-Site DC (ACI-to-ACI Intersite) Connectivity Using SD-WAN</i>	106
	Guidelines and Limitations Around ACI–SD-WAN Integration	111
	ACI and MPLS/SR-MPLS Pairwise Integration	111
	Guidelines and Limitations Around ACI MPLS/SR-MPLS Integration	132
	ACI and Public Cloud Integration	132
	Underlay and Overlay Intersite Connectivity	135
	Traffic Flows	138
	Design Option	146
	<i>Applications Stretched Across Sites (Intra-Tenant/Intra-VRF)</i>	146
	<i>Applications Stretched Across Sites (Inter-Tenant/Inter-VRF)</i>	148
	<i>Cloud to Internet (Direct Connectivity)</i>	150

*Cloud to Internet/External Destinations Using On-Prem L3Out
Connectivity* 151

Cloud to External Enterprise Sites 154

Cloud to External Enterprise Sites Using SD-WAN 157

Guidelines and Limitations Around ACI and Public-Cloud Integration 161

Summary 162

Chapter 5 Enterprise MPLS 163

Overview 163

Interaction with the Other Domains 166

Inter-working with SD-Access 166

MPLS and SD-Access Default Integration 167

Using SXP to Distribute SGTs 169

Inter-working with SD-WAN 170

MPLS as a WAN Transport 171

MPLS as a Super Core 172

Inter-working with ACI 178

MPLS with Back-to-Back VRFs 178

MPLS as a Super Core 180

Inter-working with CNF 181

MPLS to Offload Selective VRFs 181

MPLS for Inter-VRF Routing and Merging Services 182

Using MPLS to Connect NNI 184

Inter-working with Security Stack 187

Summary 189

Chapter 6 Carrier Neutral Facilities 191

Overview 191

Interaction with Other Domains 192

Inter-working with SD-Access and Campus Fabric 193

Approach with VRF-Lite 194

Dedicated Borders at the Data Center 196

Inter-working with SD-WAN, MPLS, and DMVPN 199

Software-Defined Wide Area Network (SD-WAN) 200

Dynamic Multi-point Virtual Private Network (DMVPN) 202

Provider-Multi-Protocol Label Switching (P-MPLS) 206

Enterprise Multi-Protocol Label Switching (E-MPLS) 209

Inter-working with Data Centers 214

Connecting Data Centers and CNFs with Dark Fiber	217
Connecting Data Centers and CNFs with a WAN Solution	220
Inter-working with Cloud	220
Inter-working with the Security Stack	224
Summary	227

Chapter 7 Cloud 229

Overview	229
Private Cloud	231
Public Cloud	232
Hybrid Cloud	232
Integration with Campus Networks	233
Campus Integration with Private Cloud	233
Campus Integration with Public Cloud	235
Integration with WAN	236
SD-WAN Cloud OnRamp	238
Modern Transit Network Using SD-WAN Cloud OnRamp	239
Integration with CNFs	243
CNFs and Private Cloud	243
CNFs and Public Cloud	245
Integration with MPLS	246
Integration with Security Stack	247
Summary	251

Chapter 8 Security 253

Overview	253
Security Policy	254
Security and SDA	254
Security at the SDA User Edge	256
Security Inside the SDA Fabric	257
Security and SD-WAN	262
Securing the SD-WAN Underlay	262
Securing the SD-WAN Overlay	262
SD-WAN Security and the Cloud	267
Security and DMVPN	267
Securing the DMVPN Front Door	267
Security and DMVPN	269
Security and ACI	270

Security and the Cloud	271
Security and Zero Trust	272
Summary	272

Chapter 9 Automation 273

Overview	273
CI/CD Pipeline	273
Automation and SDA	276
Automation and SD-WAN	278
Automation and DMVPN	281
Automation and ACI	282
Automation Across Multiple Domains	284
Summary	284

Part III Real-World Use Cases

Chapter 10 Manufacturing Use Case with SDA, SD-WAN, and CNF 285

Use Case Overview	285
Summary of Requirements	286
Business Requirements and Solution Mapping	287
Technical Requirements and Solution Mapping	288
Deployed Solution	290
SD-Access	290
<i>Access Layer Architecture</i>	290
<i>Core and Distribution Layer Architecture</i>	293
<i>Wireless Architecture</i>	295
SD-WAN	297
Carrier Neutral Facilities (CNFs)	298
Automation Strategy	301
Security	301
<i>Perimeter, IT, and OT DMZ</i>	302
<i>NAC and Microsegmentation</i>	304
Summary	305

Chapter 11 Financial Use Case 307

Use Case Overview	307
Campus Design Modularity	308
Campus Macrosegmentation	310
Campus Microsegmentation	312
Extending Campus Segmentation	315

SDA Data Center Services	318
SD-WAN Data Center Services	321
SD-WAN Centralized Policy	322
Summary	326

Chapter 12 Retail Use Case Using CNF, SD-WAN, and ACI 327

Use Case Overview	327
Overall Design	328
Data Center and CNF Design	330
Internet Edge	333
MPLS Edge	333
Untrusted FWs	333
SD-WAN Headends	334
Trusted FWs	335
Core	336
Cloud Edge	337
MGMT	338
CNF Backbone	338
ACI	340
SD-WAN Design	343
Advertising Routes to Remote Sites	348
Receiving Routes from Remote Sites	348
SD-WAN VPN	350
Application-Aware Policies	350
Securing Internet Access with SIG	353
Campus/Branch Design	355
Backstage Site	357
Retail Site	359
Cisco TrustSec and Meraki Adaptive Policy	359
East-West Enforcement Between Endpoints in the Same Site	362
East-West Enforcement Between Endpoints in Different Sites	367
SGT Mapping for DC Endpoints for North-South Flows and SIG Traffic	369
Hybrid Cloud Integration	371
Azure Service Deployment and Connectivity Considerations	375
Cloud Services (PaaS) Integration with ACI	378
Summary	380

Chapter 13 Public Sector Use Case 381

Use Case Overview 381

Overall Design 382

ACI MPLS-SR Integration 388

ACI Logical Design and VMM Integration 401

VMM Integration 405

Infrastructure Automation and Orchestration 413

Summary 417

Chapter 14 Transportation Use Case 419

Use Case Overview 419

Overall Design 419

Data Center Design 422

Campus Design 425

Transportation-Specific Entities 426

Campus Macrosegmentation 429

DC and WAN Macrosegmentation 431

Microsegmentation 432

SD-Access Data Center Services 433

SD-WAN Centralized Policy 437

Summary 440

Index 441

Icons Used in This Book



Border Node - Switch



Control Plane Node - Switch



Border Node - Router



Control Plane Node - Router




Edge Node




Extended Node




Transit Control
Plane Node (Switch)



Transit Control
Plane Node - Router



Fabric WLC



Fabric WLC HA SSO



Border Node Switch Stack



Border Node and Control
Plane Node Switch Stack



Edge Node Switch Stack



Colocated Border Node and Control
Plane Node with Layer 2 Handoff



Cisco DNA Center





Cisco Identity
Services Engine



Cisco vManage



Cisco vBond





Layer 3 Switch



Layer 2 Switch



Layer 3 Switch Stack



Layer 2 Switch Stack



Access Point



Access Point



Wireless LAN Controller



Wireless HA SSO



User



User Group



Laptop



Servers



Machinery



Machinery



IP Phone





















Mobile Phone



Security Camera



Credit Card Reader

 Smart Waste	 Router
 Badge Reader	 Firewall
 Smart City	 Services Block Switch
 IT Professional	 Internet
 OT Professional	 User
 Medical Devices	 Building Management Systems
 Traffic Light	 Lights
 Medical Devices	 Security Group Tag
 Street Light	 WAN Edge

Command Syntax Conventions

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

- **Boldface** indicates commands and keywords that are entered literally as shown. In actual configuration examples and output (not general command syntax), boldface indicates commands that are manually input by the user (such as a **show** command).
- *Italic* indicates arguments for which you supply actual values.
- Vertical bars (|) separate alternative, mutually exclusive elements.
- Square brackets ([]) indicate an optional element.
- Braces ({ }) indicate a required choice.
- Braces within brackets ({ []}) indicate a required choice within an optional element.

Introduction

In most enterprise organizations, different teams of architects and engineers design, manage, and operate the various domains of the environment, the wide area network, the data center, the remote campus, or the cloud environment. With larger organizations, interaction between these groups is fundamental for a well-designed global environment. The advent of intent-based networking (IBN) has given rise to policy-driven environments. The architects create a policy designed for the single specific domain managed by a given controller. Each environment has its own controller. With each domain having its own network team, its own controller, and its own policy, a uniform end-to-end architecture becomes muddled and difficult to design, operate, and manage. This book focuses on design integration in a multi-domain environment to provide a single, consistent design across the large-scale environment.

Goals & Methods

Cisco is leading the way in defining emerging networking technologies. Over the past few years, three main technologies emerged in the software-defined world: Cisco's Software Defined Access (SD-Access), Software-Defined Wide Area Networking (SD-WAN), and Application Centric Infrastructure (ACI). These technologies have been evolving ever since and gaining huge market share in consumer networks as a replacement of traditional networks. With this new way of doing things, a lot of underlying concepts are changed, and with those changes over the past few years, we are seeing customers adopting these technologies one at a time.

Many books from Cisco Press describe these technologies as an individual entity (single-domain): SD-Access, SD-WAN, ACI, Multi-Protocol Label Switching (MPLS), and so on. These are good books to build a solid foundation when starting a journey to adopt these technologies. But when the time comes to integrate these technologies within an organization to combine and interconnect those different networks, there are no guidelines, lessons learned, or best practices available for reference. Here is where this book takes the center stage.

The authors of this book—Dhrumil, Jeremy, and Navin—have more than 35 years of combined experience in building customer networks. The objective of this book is to utilize that experience so that an engineer or architect will know how to integrate these technologies and translate the organization's business and technical intent into a fully functioning, secure, and efficient network.

The technologies described here are all new and emerging, and their adoption will only increase from here in today's world of digitization and the hybrid workspace. Although these technologies themselves will continue to evolve over time, the focus of this book is on how they interact with each other rather than a deep investigation specific to one of them. Thus, evolution or revision of any one of these technologies will not impact this book.

Part I in this book starts with a chapter on answering two big questions about multi-domain: what and why? This chapter highlights the real-world demands that are driving organizations to this inevitable networking change and how they tackle adoption of multi-domain design and technologies.

Each chapter in Part II covers a high-level overview and functionality of the technology described. They describe how this technology interacts with the other technologies and identify the key items an architect would need to keep in mind while interconnecting the other technologies.

Part III spans the last five chapters of the book, covering the real-world use cases of how these technologies will seamlessly work with each other. These chapters also give some reference architecture for individuals who would want to use these technologies in tandem with each other.

The objective of this book is to give lessons learned from actual real-world deployment of Cisco's customers that the authors have captured over several years in the industry.

Who Should Read This Book?

This book is designed for architects and engineers who want to integrate these emerging technologies within their organization. The last five use cases in this book can also act as reference materials to leaders and management who are willing to make key decisions on their organization's network transformation journey.

This book is designed to be picked up from any chapter and is not meant to be read sequentially.

How This Book Is Organized

This text is not designed to be an introduction to the various Cisco intent-based networking architectures nor a detailed discussion of all of the features and functionality available in any one specific architecture. Other publications adequately handle the presentation of that material. The authors are industry experts in designing and deploying the various IBN and cloud architectures and technologies together in various permutations. This text presents valuable insights and lessons learned from previous large-scale deployments to help with planning and design.

Although you may read the entire book from cover to cover, the text was designed to be flexible, with each chapter focusing on how one specific architecture integrates with the other architectures and technologies. This allows you to combine sections from a single chapter or multiple chapters based on your solution requirements. The final five chapters of the book focus on actual industry use cases.

Book Structure

The book is organized into three parts.

Part I: Introduction

Chapter 1, “Multi-Domain Networks,” introduces the concepts of multi-domain networking. It describes how the networks have evolved to require support for multi-tenancy and reduced fault domains.

Part II: Multi-Domain Design

Chapter 2, “SD-Access and Campus Fabric,” focuses on the design and deployment of Cisco’s Software Defined Access solution while integrating it with other architectures, such as Application Centric Infrastructure.

Chapter 3, “SD-WAN and DMVPN,” illustrates the design and deployment of the various architectures with either Cisco’s Catalyst SD-WAN solution or with DMVPN.

Chapter 4, “Application Centric Infrastructure (ACI)—Integration and Multi-Domain Capabilities,” provides an in-depth discussion of the Cisco Application Centric Infrastructure integrated with various WAN, cloud, and campus technologies.

Chapter 5, “Enterprise MPLS,” focuses on the design and deployment of the enterprise MPLS environment while integrating with support for the multi-domain network.

Chapter 6, “Carrier Neutral Facilities,” describes what a Carrier Neutral Facility is, why it is critical in today’s networks, and how it integrates with other domains in the network.

Chapter 7, “Cloud,” focuses on integrating multi-domain networks with the cloud.

Chapter 8, “Security,” discusses some of the relevant security concerns and considerations when designing and deploying multi-domain environments.

Chapter 9, “Automation,” provides a discussion on extending automation and CI/CD pipelines to automate required changes in multiple domains based on the evolution of the enterprise environment.

Part III: Real-World Use Cases

Chapter 10, “Manufacturing Use Case with SDA, SD-WAN, and CNF,” illustrates a real-world deployment of SDA, SD-WAN, and CNFs integrated together.

Chapter 11, “Financial Use Case,” discusses the design requirements, caveats, and end-to-end design and deployment of a multi-domain financial sector deployment incorporating SDA, SD-WAN, and security together.

Chapter 12, “Retail Use Case Using CNF, SD-WAN, and ACI,” relies on an actual retail deployment to describe the design and deployment of SDA and SD-WAN with a cloud environment.

Chapter 13, “Public Sector Use Case,” uses a real-world scenario to illustrate SDA and ACI integration with an enterprise MPLS environment.

Chapter 14, “Transportation Use Case,” illustrates the end-to-end design and integration of SDA, SD-WAN, and ACI to address multi-tenancy in transportation.

Figure Credit

Figure 4.1a: renjithkrishna/Shutterstock

This page intentionally left blank

SD-WAN and DMVPN

In this chapter, we discuss the following:

- Common design strategies in SD-WAN and DMVPN deployments
- How to design and deploy SD-WAN to integrate with other IBN domains
- How to design and deploy DMVPN to integrate with other IBN domains

Overview

Both Cisco Software-Defined WAN (SD-WAN) and Dynamic Multi-point Virtual Private Network (DMVPN) provide the ability to abstract the WAN service provider transports from the enterprise routing environment. Additionally, both provide a means to create and extend macro- and microsegmentation, including support for Cisco TrustSec. This support allows either architecture to be utilized as part of an end-to-end security policy. Cisco SD-WAN has many advantages as an architecture over DMVPN, such as application-aware routing and built-in automation and provisioning; however, DMVPN does have its use cases. Both of these technologies fundamentally provide an efficient way of routing between the sites by providing direct site-to-site communication without the need for going through a centralized hub or a data center.

SD-WAN

Cisco SD-WAN as a technology is discussed in detail in various other Cisco Press texts. The sections in this chapter assume that you are familiar with Cisco SD-WAN. The following sections discuss designing and integrating Cisco SD-WAN with the various other domains as part of a single multi-domain strategy.

SD-WAN and SDA

Cisco's Software Defined Access, or SDA, allows the enterprise to introduce macro- and microsegmentation with automation and assurance at the local campus environment. Hosts may be dynamically or statically classified into virtual networks (macrosegmentation) and security groups (microsegmentation). When discussing SDA, it is important to remember that the control plane is based on LISP, whereas the data plane uses VXLAN.

In a multi-site SDA deployment without integration into other domains, the architect must plan for policy enforcement and ensure that the virtual network identifier (VNID) and security group tag (SGT) information is correctly propagated across either an IP transit environment or via SDA-Transit. SD-WAN allows the enterprise to extend the macro- and microsegmentation end to end in a fully integrated fashion. This way, the SGT information can be propagated across the network without additional resource utilization or reclassification as the data re-enters the SDA environment at the remote location. Even if the remote site has not been migrated to SDA, SD-WAN may be utilized to provide Cisco TrustSec policy enforcement at the remote site without the originating site being aware of the destination SGTs.

Cisco product engineering has supported two methods for integrating SDA and SD-WAN: the one-box and the two-box solutions. One of the most essential points to remember is that inline tagging for SD-WAN is only supported on routers based on Cisco platforms such as the ISR 4000 series, ASR 1000 series, and Catalyst 8000 series. Viptela-based routing platforms, such as vEdge 100 and vEdge 1000, are *not* supported for inline tagging.

One-Box SDA and SD-WAN

One-box SDA and SD-WAN is also known as an *integrated solution*. In this scenario, the SD-WAN Edge router serves as an SDA control plane and also a border node. For this to occur, the Cisco DNA Center must be integrated with the Catalyst SD-WAN Manager via API integration. The SD-WAN Edge will be a part of the Catalyst SD-WAN Manager inventory and be provisioned as a normal SD-WAN device. When the SDA fabric is created, the Cisco DNA Center will update the Catalyst SD-WAN Manager via the API integration to reprovision the SD-WAN Edge with the appropriate SDA configuration. Figure 3-1 shows how a packet changes as it traverses the one-box solution across SD-WAN from one SDA site to another. Notice that the VNID and SGT information is propagated across the network with the data packet itself.

There are distinct pros and cons to the one-box approach that must be considered as part of the overall enterprise design. As will be discussed with the two-box approach, the one-box approach removes support for modularity. This is important to consider because most enterprises are divided into multiple organizational units even when considering who manages which part of the network. For instance, one group may manage the WAN while another manages the LAN campus. The one-box approach will make it difficult for the two groups to manage their respective environments.

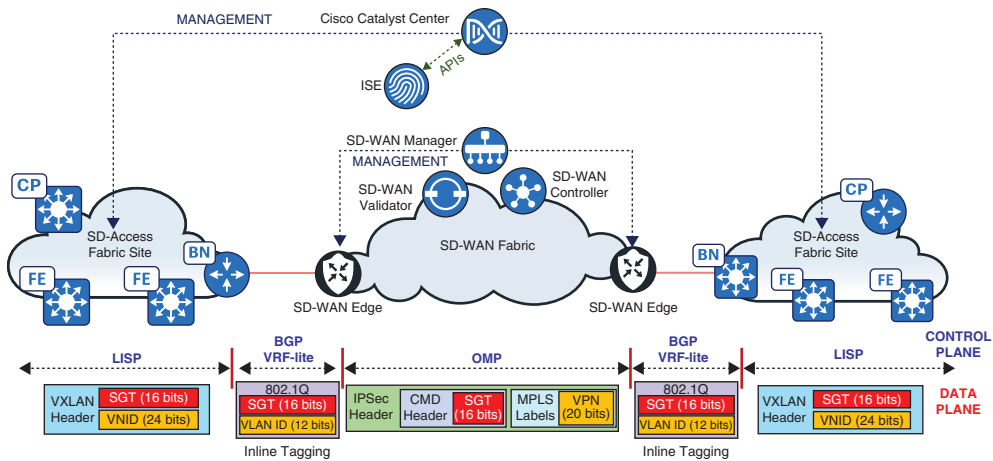


Figure 3-1 SD-WAN-SDA One-Box Topology

From a design consideration, it should be noted that the one-box solution requires a router to perform the border node and control plane functionality in the SDA campus. With physical redundancy included, two routers now perform those roles. While this solution works well for the control plane, because the routers have greater routing capabilities, in the data plane, the routers could inadvertently become the logical core of the local area network. In larger locations, additional functional blocks may exist outside of the SDA domain. For instance, where Nexus switches perform a local services aggregation block, an additional pair of switches performing the internal border node functionality at the intended core will better facilitate the required high-speed switching without the traffic traversing the edge routers. Figure 3-2 illustrates how the connectivity between the additional non-SDA fabric at the local site could connect directly to the core of the network while still utilizing the one-box solution. Notice that the core now performs SDA border node functionality. SDA VXLAN traffic that is egressing the location will still utilize the SD-WAN Edges, whereas traffic to these additional services will utilize the core border nodes as their VXLAN termination point.




Figure 3-2 One-Box Topology with Additional Service Domains

Catalyst SD-WAN Manager and Catalyst Center Integration

Integrating Catalyst SD-WAN Manager and Catalyst Center is fairly straightforward. From the Cisco Catalyst Center System Settings, navigate to the External Services page and select **Catalyst SD-WAN Manager**. Depending on the version of Catalyst Center, this selection will navigate to a new page or open a pop-out, allowing you to enter the Catalyst SD-WAN Manager and SD-WAN overlay information shown in Figure 3-3.

The Catalyst Center will use the configured user credentials for its API calls to Catalyst SD-WAN Manager. Therefore, it is recommended that a service account for the integration is created within the enterprise identity store that can be used for auditing purposes. As noted in Figure 3-3, if the Catalyst SD-WAN Manager is authenticated via a root certificate authority, then the Catalyst Center must have a certificate installed through the Certificates page from the same trust chain.





Host Name/IP Address*	172.31.23.236	
The hostname or IP address of vManage		
Username*	admin	
The user ID of vManage		
Password*	*****	 SHOW
The password of vManage		
Port Number*	8443	
The vManage port number		
vBond Host Name/IP Address	vBondhost	 Info
Organization Name		
	sdwan-overlay	 Info

Figure 3-3 Catalyst SD-WAN Manager and Catalyst Center Integration Process

Two-Box SDA and SD-WAN

In the two-box SDA–SD-WAN scenario, also known as a *nonintegrated solution*, both architectures are kept separate, providing for a modular networking approach. The SD-WAN Edge devices each have a physical link to the SDA border nodes that is a dot1Q trunk. On the SD-WAN Edge side of the link, the subinterface is associated with a

specific service VPN. On the border node side of the link, the SVI provisioned from Catalyst Center is associated with the corresponding SDA virtual network. In this manner, the VLAN becomes the mapping between the SDA virtual network and the SD-WAN service VPN. By enabling Cisco TrustSec (CTS) inline tagging on both sides of the interface, the SGT value may be propagated as well, maintaining both the macro- and microsegmentation of the environment.

From a design perspective, the two-box scenario allows for both a modular design and phased rollouts at the expense of potentially more hardware to install and manage. Figure 3-4 demonstrates how a data packet traverses the two-box solution while maintaining the VNID and the SGT information with the packet itself.

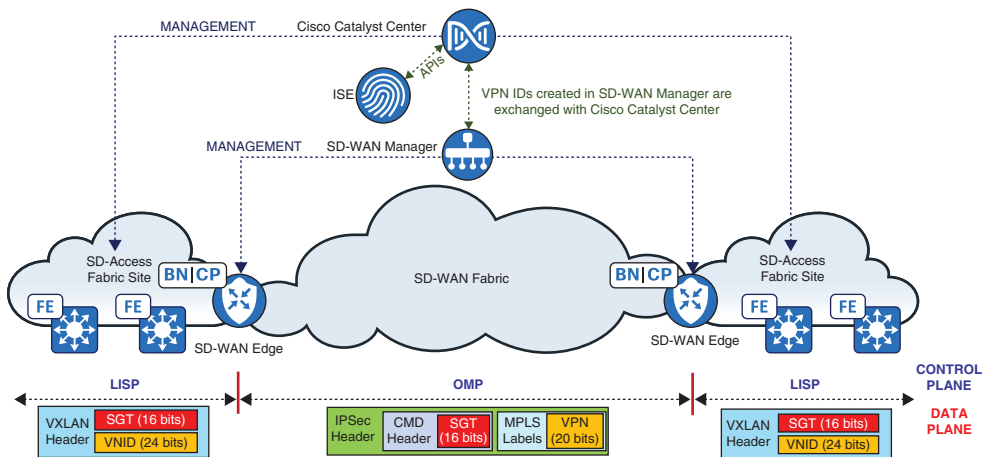


Figure 3-4 SD-WAN-SDA Two-Box Topology

SDA and SD-WAN Segmentation

In SDA, the VNID is a 24-bit value used to identify which virtual network a packet traversing the underlay is associated with. When a new virtual network is created in the Catalyst Center UI, Catalyst Center creates a new VNID for it that is constant across all locations. Whenever that virtual network is provisioned at an SDA site, Catalyst Center uses the VNID as the LISP instance ID, which is mapped to the VRF on the switch with the correct virtual network name. The VNID is carried across the underlay as part of the VXLAN header. Additionally, the VXLAN header carries the SGT value. The SGT is a 16-bit value that indicates to which security group the source of the packet belongs.

For data egressing the SDA environment, it is forwarded as a VXLAN packet to the correct border node. After the packet arrives at the border node, it is decapsulated and forwarded based on the VRF instance associated with the VNID. If the destination security group is known by the border node, it will enforce applicable policy. This is not always the case because it would depend on the configuration of the border node.

With SD-WAN, the service VPN ID is inside the IPsec header prior to forwarding on the transport layer. Additionally, the CMD header commonly associated with Layer 2 frames has been added to the IPsec header on an SD-WAN packet to allow the SGT information to be propagated also.

The last piece of the discussion pertains to connecting the SDA environment and the SD-WAN environment together. Whether the one-box or two-box solution is utilized, there must be a consistent mapping between the two architectures. In the one-box solution, mapping is done via Cisco Catalyst Center. After the Catalyst SD-WAN Manager to Catalyst Center integration has been performed as described previously, the individual SDA virtual networks are mapped to specific SD-WAN service VPNs in the Catalyst Center UI. When an SD-WAN Edge is provisioned at an SDA location, the mapping of Service VPN to VNID configured in the Catalyst Center is used. The SD-WAN Service VPN is used as the name of the VRF on the SD-WAN Edge for all of the SDA and SD-WAN appropriate configurations. When the integration is complete, the Catalyst Center page utilized to tie the SD-WAN Service VPN to the SDA virtual network is shown, as in Figure 3-5.

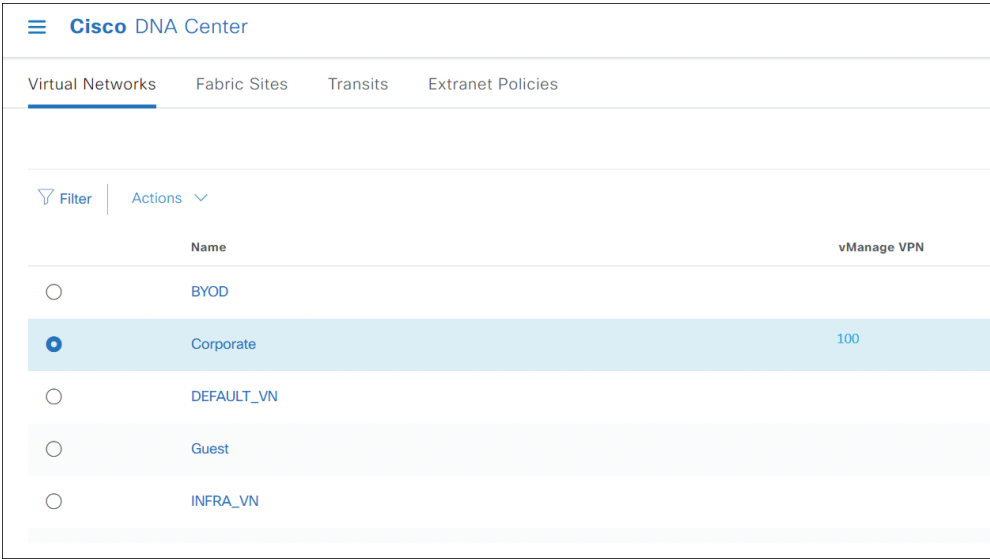


Figure 3-5 One-Box VNID to Service VPN Mapping

In the two-box solution, the mapping between VNID and service VPN is achieved through the use of the VLAN carrying the traffic between the two devices. It is critical to standardize the VLAN mapping across the environment. Doing so facilitates easier operation and troubleshooting of a multi-site environment when operations engineers know that a specific VLAN is used to connect SD-WAN Edge to the border node in the SDA Corporate VN and SD-WAN service VPN 500 at all locations. The VLAN mapping ensures that macrosegmentation is maintained; however, support for the microsegmentation propagation must be intentionally added. This is achieved by configuring inline

tagging on both sides of the link, allowing the SGT information to be propagated via the CMD header in the frame. Care should be taken to ensure that the device SGT also is trusted on both sides. In Example 3-1, the inline tagging is configured on both the physical interface and the subinterface carrying the traffic. The SGT with ID number 2 is the All Cisco TrustSec Devices SGT.

Example 3-1 *Inline Tagging Configuration*

```
! LAN Interfaces
interface GigabitEthernet0/0/0
  cts manual
  policy static sgt 2 trusted
!
interface GigabitEthernet0/0/0.100
  cts manual
  policy static sgt 2 trusted
!
```

SDA and SD-WAN Best Practices

As is seen throughout this book, standardization across the environment is critical. When creating any standardized mapping, ensure that future proofing is considered. Is there a potential for additional network devices for horizontal scaling at the location? Will additional macrosegmentation be added to the environment that may need to be considered?

The SDA INFRA_VN, or SDA underlay, should be contiguous with a service VPN in SD-WAN. The SDA underlay may be mapped to a unique SD-WAN service VPN or utilize one used by the general corporate VPN. The former allows the underlay management to be reachable from the local environment without traffic leaving the site, and the latter maintains the macrosegmentation requiring traffic to egress the site, possibly fusing the two routing domains together at a centralized firewall environment.

When the SDA and SD-WAN multi-domain environment is built out, it is recommended to build the centralized data center environment first—the services and SD-WAN Edge headend devices. At the remote sites, the process of building and validating the SD-WAN environment first facilitates a smoother transition to deploying SDA.

When Cisco TrustSec is enabled at a site, care should be taken to ensure one device does not become inaccessible. In a location where physical redundancy exists, the engineer should ensure that the SDA devices are accessible across both pathways prior to enabling CTS. Additionally, it is important to remember that CTS must be enabled on the physical and subinterfaces on the link.

In the event that it is a small site with minimal hardware and connectivity—for instance, a single SD-WAN Edge cabled to a single Fabric-in-a-Box (FiaB) switch—redundant pathways may not exist. In this instance, it is recommended to have the CTS configuration as a simple text file on the flash of the switch. The file may be copied into the running

configuration on the switch. At this point, the reachability to the switch will be lost; however, all of the configuration in the text file will be applied. The SD-WAN device may then be reprovisioned to include the CTS, restoring the connectivity.

SD-WAN and ACI

Cisco's Application Centric Infrastructure (ACI) allows the enterprise to introduce macro- and microsegmentation with automation and assurance within the data center. ACI uses a structured hierarchy including tenants, contexts, and endpoint groups (EPGs) to create macrosegmentation and microsegmentation. The EPG is similar to the SGT in the SDA and SD-WAN environments. It allows for policy enforcement based on logical group membership.

When SD-WAN and ACI are deployed, the individual macrosegmentation that is created within the DC ACI environment is extended to the remote site location. For instance, the enterprise may provide managed services to their end customers, internal or external, and want to ensure segmentation from the DC to the site. Having an SD-WAN service VPN for each ACI tenant will maintain that segmentation. While the current APIC and Catalyst SD-WAN Manager allow for integration via REST APIs, that integration only supports a dynamic application-aware routing policy signaling from ACI to SD-WAN. Therefore, all of the routing interconnectivity must be performed individually in both environments. For this reason, standardization again becomes important.

Imagine an enterprise environment without SD-WAN or ACI consisting of two data centers and multiple remote locations. This enterprise currently has all of its clients in a single global routing table without any segmentation. Now, they would like to migrate to full segmentation with both ACI and SD-WAN deployed. It will take some time to build the ACI environment and migrate the relevant services into each tenant. It will also take time to migrate each of the individual sites to SD-WAN. How is this performed without issues?

First, we must consider all of the possible traffic patterns. There is traffic from the nonmigrated data center environment to the nonmigrated remote locations through the service provider environment using the current CE equipment. This traffic will exist throughout all the migrations until both the SD-WAN and the ACI migrations are fully completed, although the amount of traffic will decrease with each migration window. As the migrations proceed, there will be traffic from the ACI environment through the SD-WAN environment to the remote locations. At first, this traffic will not exist at all and will increase as migrations occur. There will also be traffic between the nonmigrated data center environment and the new SD-WAN remote locations, as well as nonmigrated remote sites with the newly migrated ACI environment. Additionally, traffic will exist between migrated and nonmigrated sites, and an existing data center with the ACI environment.

All of these traffic patterns will exist in some amount from the beginning of the project until the end. Therefore, from a routing and switching perspective, there are four domains: the SD-WAN environment, the ACI environment, the existing data center, and

the existing WAN environment. It is recommended to create an additional routing and switching layer within the data center that performs aggregation and routing between the domains. In Figure 3-6 notice a new aggregation layer has been inserted between the legacy WAN environment, the new SD-WAN devices, the legacy data center services infrastructure, and the new ACI environment. This new layer will allow the routing to drive traffic to the correct blocks based on the destination location—whether already migrated to the new environment or not.

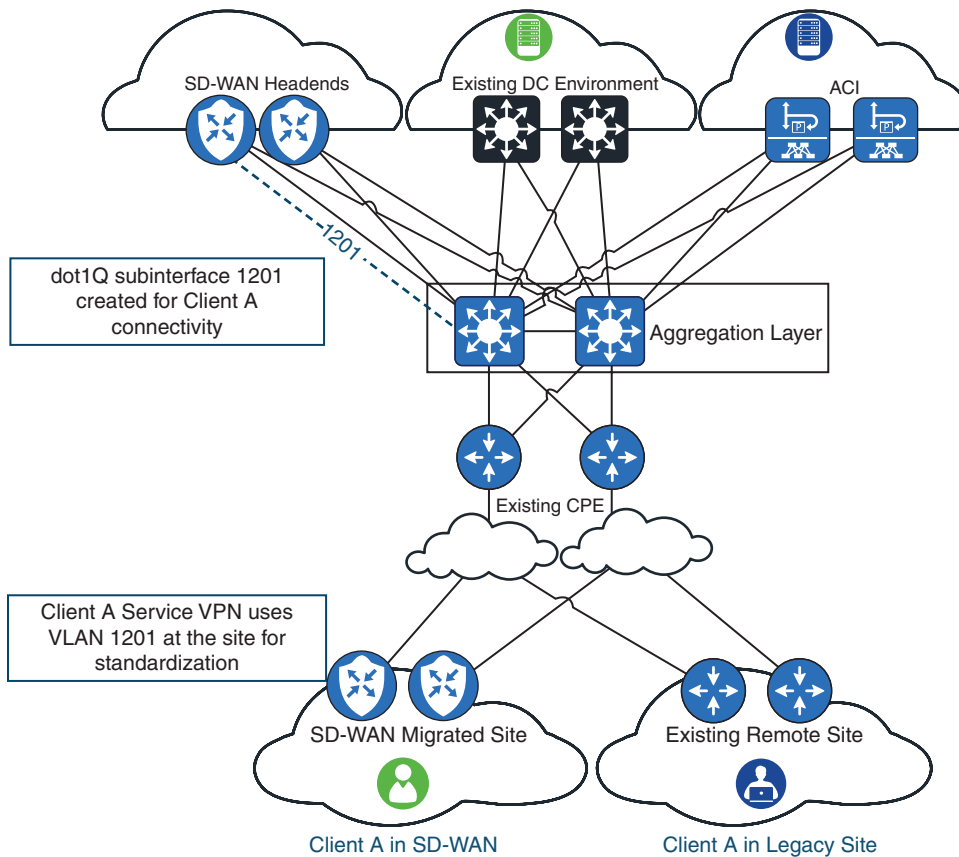


Figure 3-6 SD-WAN-ACI Topology

When the environment is designed and implemented, the use of standardized VLANs will facilitate an easier migration per client. After an aggregation layer is created within the data center to facilitate interactions between the environments, the SD-WAN headend devices may be stood up appropriately. When this has happened, the ACI and SD-WAN environments are migrated at their own individual rates. This approach allows the WAN team to focus on just the remote location migrations while the data center team is able to focus on client services.

Consider the migration of Client A. At first, the services for the client exist in the existing data center environment, and the remote locations that service this client all use the global routing table with the service provider. When the aggregation layer and the SD-WAN headends are in place, the migration of the client is transparent to the client, with the exception of the required routing updates during maintenance windows. Perhaps the ACI environment is not built out while the SD-WAN environment is ready for production. The service VPN for this client is provisioned on the headends—for example, VPN 1201. As part of the provisioning, BGP peering between the headends and the aggregation layer on VLAN 1201 is created. Provisioning the new service VPN in the headends will have no effect on the traffic flows because there will be no routing advertisements at this point coming from the headends. Whenever a remote location is moved to VPN 1201, the headends will begin to advertise the remote site via BGP while the service provider will lose the routing advertisement from the remote location. This is the case with Client A.

At any specific remote location, there may be a different collection of service VPNs, that is, clients, from other locations. Because it is conceivable that the headend environment may not be provisioned for all clients or the enterprise wants to move to SD-WAN quickly, we may want to create a single-service VPN that may be used similarly to the existing function of the global routing table. That is, migration to SD-WAN is performed, but segmentation is not fully introduced. Perhaps the local network has not been configured for segmentation via VRF Lite or some other manner; this common service VPN allows for the entire site to move to SD-WAN while not affecting the local environment. When the local network is ready to migrate Client A to its own segmented environment, the Client A service VPN is provisioned on the SD-WAN Edges at the site. With the ensuing routing updates, the Client A traffic for this remote site now uses the SD-WAN environment and is advertised from the headends to the data center aggregation layer to all other environments.

This architectural design also works for the migration of services for Client A. When the ACI environment is ready for production, all of the logical ACI components are added into the ACI environment to support Client A. The required services for Client A are moved into the ACI environment, and the ACI L3Outs, or border leafs, advertise the services to the data center aggregation layer.

Therefore, while ACI and SD-WAN are integrated together, the migration of the services for a particular tenant in ACI and the migration of the tenant's remote locations in SD-WAN may proceed at their own individual pace. The aggregation layer handles the routing between the various environments. As shown in Figure 3-7, the aggregation layer allows a remote site that has been migrated to SD-WAN already to interact with a remote site that has not. The reason is that the SD-WAN headends are advertising to the aggregation layer the SD-WAN remote site prefixes to the legacy WAN environment, and vice versa. With an L3 MPLS offering from the service provider, it is conceivable that these sites are able to send traffic directly to each other across the service provider. However, because the SD-WAN traffic is encrypted while the legacy traffic across the service

provider is not, during this hybrid state of migrated and nonmigrated sites, the headends and the aggregation layer must be utilized to interconnect the domains.

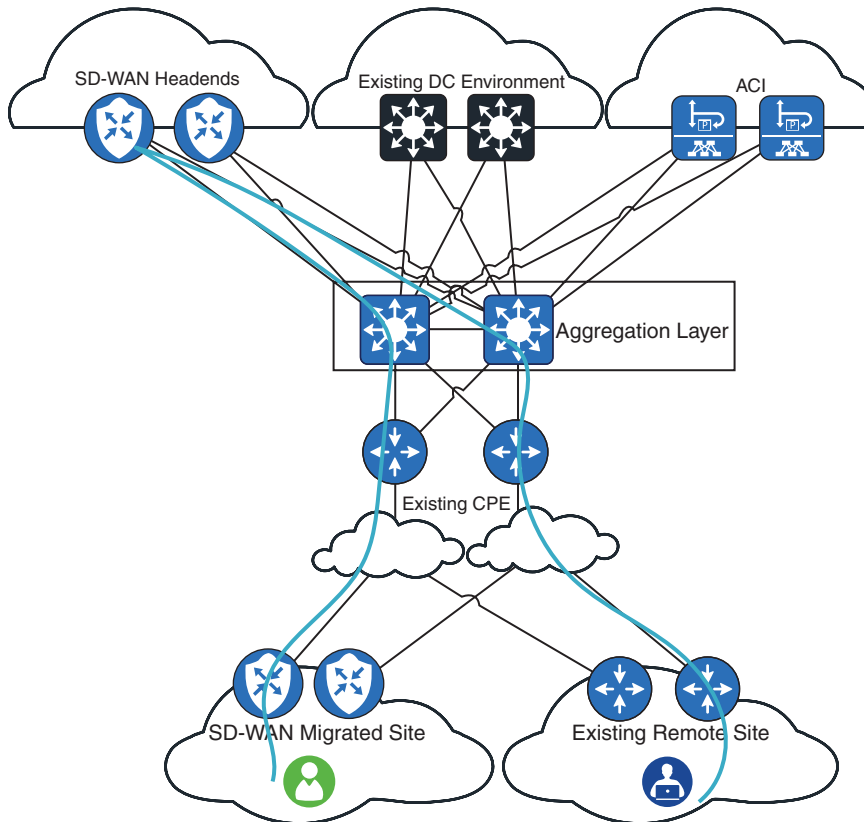


Figure 3-7 SD-WAN-ACI Traffic Flows During Migration

Catalyst SD-WAN Manager and APIC Integration

Integration of the Catalyst SD-WAN Manager with the ACI APIC is performed on the APIC itself. A static user on the Catalyst SD-WAN Manager is required for the APIC to communicate with the Catalyst SD-WAN Manager. For security, auditing, and best-practice purposes, it is recommended to use a service account for the integration, as well as authentication via an external identity store. Doing so will facilitate proper user auditing, as well as the ability to manage the account via the appropriate operations processes.

Example 3-2 illustrates the configuration process required to integrate the APIC and the Catalyst SD-WAN Manager together.

Example 3-2 *Catalyst SD-WAN Manager and APIC Integration Process*

```

apic1#conf t
apic1(config)#integrations-group MyExtDevGroupClassic
apic1(config-integrations-group)#integrations-mgr External_Device Cisco/vManage
apic1(config-integrations-mgr)#device-address 172.31.209.198
apic1(config-integrations-mgr)#user admin
Password:
Retype password:
apic1(config-integrations-mgr)#

```

ACI and SD-WAN Segmentation

While ACI and SD-WAN both support the concepts of macro- and microsegmentation, microsegmentation propagation does not occur without additional configuration. Also, the macrosegmentation propagation must be handled in a systematic manner.

For macrosegmentation propagation, this is where the ACI Tenant to VLAN to Service VPN mapping is important. The VLAN used to interconnect the ACI L3Out, or border leaf, to the SD-WAN Edge is crucial to maintain the macrosegmentation.

Microsegmentation propagation of the ACI EPG to SD-WAN SGT values is more difficult and limited. The APIC must be integrated with ISE using pxGrid in the same manner as used for the ACI-SDA integration. This will allow ACI to advertise or receive EPGs to and from ISE; however, as with the ACI-SDA integration, this is limited to a single context. For the SD-WAN side, the headend SD-WAN Edges may use SXP with ISE in any of the service VPNs. The SGT information will be propagated along the data path of SD-WAN to the remote SD-WAN Edge, where policy enforcement or further propagation may occur.

ACI and SD-WAN Best Practices

For ACI and SD-WAN integration together, the two most important aspects are standardization of the handoff between them and the support for migrated and nonmigrated traffic flows. For the former, it is recommended to use a planned VLAN to Service VPN numbering. This approach prevents confusion later when some clients have migrated to ACI while others have not, or when some sites or clients have been migrated to SD-WAN while others have not. For the latter, the use of the aggregation layer with BGP allows the enterprise to connect the legacy and new environments together while using BGP to affect policy routing, if necessary.

SD-WAN with MPLS

Because SD-WAN is designed to utilize multiple independent transports from various service providers, integrating SD-WAN with an MPLS deployment is rather straightforward.

There are numerous reasons why the enterprise may want an SD-WAN deployment while still utilizing an L2VPN or L3VPN MPLS deployment.

The first scenario is quite obvious: migration from MPLS to SD-WAN. In this scenario, the enterprise already maintains a WAN topology facilitated by their MPLS provider and plan to move to SD-WAN, perhaps to take advantage of less expensive broadband circuits. However, there are other scenarios where both the MPLS environment and the SD-WAN environment will coexist by design. For instance, the interconnection between data centers may be through an L2VPN offering that should be maintained even after the remaining WAN has migrated or deployed SD-WAN.

In most scenarios, the primary concern for design and implementation will be on proper route filtering. Depending on the routing protocols utilized within the environment, it is possible to inadvertently cause a routing loop through redistribution, as well as create suboptimal routing. For this reason, it is recommended that all best practices around route redistribution are strictly followed, including marking all prefixes that are redistributed from one protocol to another. This may be via OSPF and OMP tags, BGP communities, and so on.

SD-WAN has various built-in mechanisms to prevent route looping. For instance, when the OMP overlay AS has been configured, this ASN is added to the BGP as-path attribute when the SD-WAN Edge advertised the prefix into BGP. Additionally, when the SD-WAN Edge advertises an OMP prefix into OSPF, the down bit is set. The SD-WAN Edge works in a similar fashion to an MPLS PE node. However, without proper care, the mechanism used by SD-WAN to prevent route looping could be bypassed.

SD-WAN and the Cloud

Over the past several years, enterprises have started to move heavily into the cloud. This is true with many of the largest enterprises that had been traditionally cloud averse. The strong push to a hybrid work environment, as well as the availability of integrating with Secure Access Service Edge (SASE) architectures, has helped facilitate the migration to the cloud. Additionally, SD-WAN itself not only participates in SASE but also offers various solutions via Cloud OnRamp to assist in deploying into the cloud.

Cisco's SD-WAN offers three virtual platforms for extending the SD-WAN environment into the cloud: the CSR1000V, the vEdge-cloud, and the Catalyst 8000v. The first two here are approaching the end of life, so the virtual platform of choice moving forward should be the Catalyst 8000v. Both Amazon Web Services (AWS) and Azure offer the Cat8kv with multiple compute options in various zones and regions. As with any cloud virtual deployment, the compute requirements should be carefully considered based on throughput requirements, as well as overall cost. For instance, there are scenarios where doubling the compute for a virtual SD-WAN Edge in AWS doubles the cost of the VM; however, the throughput of the SD-WAN Edge itself is not doubled. In this scenario, it is more cost-effective to double the number of virtual SD-WAN Edges deployed in AWS. Doing so not only doubles the cost and the compute resources but also the total amount of throughput in the virtual environment. Therefore, horizontal scaling in the cloud is not just a useful practice for applications but also for the virtual network functions.

Discussing the cloud deployment is not dissimilar from any other network deployment. Does the environment constitute a greenfield deployment or brownfield? In SD-WAN, that question is even more important than usual because the current Catalyst SD-WAN Manager versions support only greenfield integration for certain Cloud OnRamp features. Therefore, if, for instance, the deployment already has VPCs in AWS that the enterprise wants to deploy SD-WAN virtual routers into, the Catalyst SD-WAN Manager Cloud OnRamp workflows will not work in that scenario. However, whether it is brownfield or greenfield, the overall design will be the same with the differences coming from how the virtual routers are deployed and maintained.

Either way, the virtual cloud SD-WAN Edge is configured from Catalyst SD-WAN Manager via templates just like any other SD-WAN router. The cloud environment itself may be considered to be another site in the SD-WAN environment. As with all routers, there is a finite number of tunnels and throughput the virtual router may support, so the control policy should be defined to ensure those thresholds are not exceeded.

SIG

One of the fundamental pieces of SASE is Secure Internet Gateway (SIG). As applications have moved to the cloud, such as Microsoft Office 365, the traditional paradigm of Internet direct from the data center or centralized location has created bottlenecks in network throughput because the Internet circuits in the centralized location were not deployed for all of the application traffic. As such, enterprises look to offload the Internet application traffic at the remote site. However, this opens new concerns from a security perspective, especially because the data center environment is normally built with security inspection and defense in depth in mind.

How then do we secure the remote site Internet edge, ensuring that application traffic is inspected without additional hardware? The first part of the answer is SIG. With SIG, the SD-WAN Edge will use API calls to the cloud service, commonly Cisco Umbrella or other third-party vendor solutions. The API calls to the cloud service are used by the SD-WAN Edge to create a direct point-to-point encrypted tunnel to the service provider. With the addition of a SIG service route to steer Internet-destined traffic or specific traffic applications across the SIG tunnel, the remote-site application traffic specified by the policy is sent encrypted to the provider. Depending on the policy and service offering, the provider then performs the required inspection on the application traffic. The provider uses NAT Translation of the application traffic so that return traffic for the application is returned to the cloud prior to sending to the remote site over the encrypted tunnel.

As with almost all technologies in networking, SIG supports redundancy. We may configure active/standby tunnel pairs where one tunnel terminates in one zone or region, and the other tunnel in the pair terminates in another zone or region of the provider. Also, the SD-WAN solution probes across the tunnel to monitor state, so the application traffic may be steered through the data center in the event that the SIG pathway is not viable. Up to four active/standby tunnel pairs may be configured on a single SD-WAN Edge to achieve maximum throughput performance for the SIG tunnels as a single tunnel throughput is capped based on the software version.

In Figure 3-8, traffic destined to the enterprise uses the SD-WAN fabric across the various service providers following the various SD-WAN policies; however, traffic that is destined for the Internet follows the encrypted SIG tunnel to the SIG service provider.




Figure 3-8 *SD-WAN SIG Traffic*

Cloud OnRamp

The Cisco SD-WAN solution offers several enhancements as part of the Cloud OnRamp (CoR) features that facilitate SD-WAN cloud connectivity. Cloud OnRamp for SaaS allows the SD-WAN solution to integrate and properly steer application traffic for select applications that are cloud hosted, such as Office 365, Dropbox, and others. With CoR SaaS, the solution probes the pathway through the DIA circuit from the site, as well as the pathway through the data center via the normal SD-WAN tunnels. Based on the probe performance and configured policy, the SaaS application traffic is steered appropriately between the options. Cloud OnRamp for IaaS handles the provisioning of virtual SD-WAN Edge devices within the cloud provider, AWS, or Azure. As part of the provisioning of the environment, the appropriate VPCs or VNets are configured based on the workflow. Additionally, Software-Defined Cloud Interconnect (SDCI), which evolved from the Cloud OnRamp for Multicloud workflow, allows for the creation of middle-mile topologies. In these workflows, the SD-WAN Edges at remote sites create SD-WAN tunnels to one of the two supported providers, Equinix or Megaport. The provider then

provides SD-WAN tunnels direct to the cloud provider over the provider's infrastructure, reducing the requirement on Internet traversal. All of these Cloud OnRamp options may be used separately or together. This scenario is illustrated in Figure 3-9.




Figure 3-9 *SD-WAN Cloud OnRamp for SaaS*

In this figure, user application traffic destined for one of the SaaS providers uses the direct Internet access at the SD-WAN site directly. All other traffic follows the SD-WAN fabric pathways. Configuring CoR SaaS within Catalyst SD-WAN Manager is fairly straightforward. From the Administration Settings page within Catalyst SD-WAN Manager, enable Cloud OnRamp for SaaS. Additionally, Cloud Services and Catalyst SD-WAN Analytics must be enabled from the same page. This will require entry of a one-time password and cloud gateway URL that are provided at the time of system setup. After the feature is enabled, you can use the Cloud OnRamp for SaaS configuration pages to view and manage how the SaaS applications should be monitored. Additionally, support for SaaS can be systematically deployed across the environment on a per-site basis as required.

Setting up Cloud OnRamp for IaaS or Cloud OnRamp for Multicloud requires associating the cloud service provider account. As of the 20.9 Catalyst SD-WAN Manager UI, the CoR IaaS functionality is moved into the Cloud OnRamp Multicloud page. Because these are enterprise accounts, it is again recommended to follow best practices and security operations requirements around creating a service account for this part. After the

appropriate account has been configured within Catalyst SD-WAN Manager using the Associate Cloud Account workflow, the UI allows the user to associate and tag the VPCs that will then be used within the Intent Management. The Intent Management piece is where the branch-to-cloud connectivity is defined within the workflow.

The same workflows allow the user to create middle-mile connectivity through either Megaport or Equinix via the Software-Defined Cloud Interconnect controls. Just as following the workflows allows cloud SD-WAN Edges to be provisioned in AWS or Azure, these workflows allow the circuits between middle-mile locations to be allocated as required. Figure 3-10 shows the various cloud and on-premises environments that may be interconnected via SDCI.

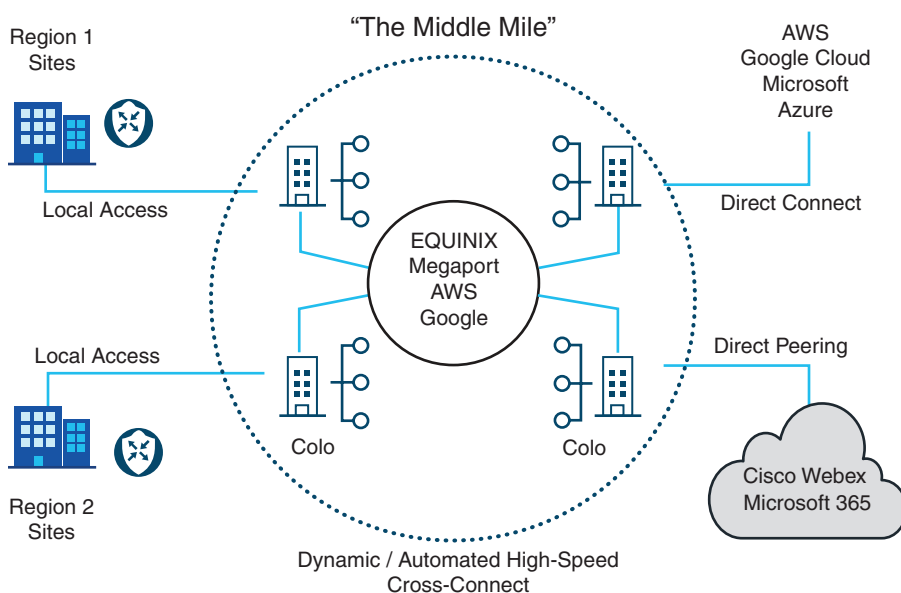


Figure 3-10 *SD-WAN Software Defined Cloud Interconnect*

As shown in the figure, with the SDCI working in the middle of the architecture, SD-WAN is capable of creating dynamic tunnels between sites and the nearest colocation facilities. The facilities themselves then provide direct peering to application providers, direct connection to other cloud services, or global connectivity to other regions and colocation facilities.

DMVPN

Dynamic Multi-point Virtual Private Network (DMVPN) is an older tunneling technology for WAN utilization that allows for simplified configuration of a hub-and-spoke topology while offering support for dynamic spoke-to-spoke tunnels. The technology also includes support for IPsec encryption across the tunnels, allowing for secure communications between all of the enterprise locations.

Similar to SD-WAN, DMVPN is capable of carrying the SGT information across the encrypted tunnel architecture. While there are numerous advantages and benefits when using SD-WAN as compared to DMVPN, numerous DMVPN deployments exist in production today.

The details of how DMVPN itself works are discussed in various Cisco Press texts. For questions on configuring DMVPN, please review those works.

DMVPN and SDA

One of the more common DMVPN multi-domain strategies is the integration with SDA. In this scenario, the enterprise has multiple locations where each is an SDA fabric site interconnected via DMVPN. Several approaches may be taken here, depending on the enterprise's end goals. Therefore, a clear understanding of the extent of macrosegmentation in the DMVPN environment is important for the design. For instance, the goal may be to extend macro- and microsegmentation throughout the DMVPN environment extending into the hub locations. Additionally, what is the design for any guest network or other SDA virtual network that is using VN anchoring? Are there locations where the local DMVPN router should behave as a fusion device between two or more SDA virtual networks?

For each SDA virtual network that requires segmentation across the DMVPN topology, a unique DMVPN tunnel in that VRF must be configured. While existing DMVPN deployments may have already used the global routing table for the DMVPN tunnel source, it is recommended to use an isolated frontdoor VRF (fVRF) for the tunnel source with the service provider. This fVRF provides isolation in the routing table information between the global or corporate routing information that traverses the tunnels and the tunnel source information with the service providers. This also inherently prevents any tunnel recursion issues that must be considered when the tunnel source and the overlay are in the same routing instance. When SDA macrosegmentation integration is added with the DMVPN environment, this will further reduce the overall complexity by having a single VRF used as the fVRF for all of the tunnels across all VRFs.

Note that the various DMVPN tunnels at the spoke sites may not be required at all locations. The tunnel requirement would be based on what SDA virtual networks are configured at a specific location. In Figure 3-11, the red virtual network is extended from one SDA location to one remote location while the blue virtual network exists at both SDA locations and the other remote location. The extension of the SDA virtual network with the DMVPN tunnel topology depends on the placement of the virtual networks, as well as the individual enterprise requirements. From the DMVPN topology to the SDA border node is a dot1Q trunk mapping VLANs to SDA virtual networks. As mentioned in other chapters throughout this book, using standardized VLAN mapping will facilitate smoother and simpler operations in production.

After a DMVPN tunnel has been configured for the respective SDA virtual network, adding SGT inline tagging to the tunnel and the interface between the SDA BN and the DMVPN router will allow for microsegmentation propagation.




Figure 3-11 *DMVPN-SDA Topology*

Example 3-3 illustrates the use of the frontdoor VRF configuration to isolate the transport while the tunnel interface itself is in the corporate VRF with the SGT value propagated. This sample configuration would be for the headend while the remote site would have the relevant NHRP configuration changes.

Example 3-3 *DMVPN-SDA Configuration*

```
! Per-VN Tunnel Interface
interface Tunnel0
  vrf forwarding Corporate
  ip address 10.0.0.1 255.255.255.0
  tunnel vrf Underlay
  no ip redirects
  ip mtu 1400
  ip nhrp authentication test
  ip nhrp map multicast dynamic
  ip nhrp network-id 100000
  ip nhrp shortcut
  ip nhrp redirect
  ip tcp adjust-mss 1360
  tunnel source GigabitEthernet0/0/0
  tunnel mode gre multipoint
  tunnel key 100000
  tunnel protection ipsec profile profile-dmvpn shared
  cts manual
  propagate sgt
!
```


DMVPN and SDA Policy Enforcement

As mentioned previously, the design requires an understanding of how DMVPN will be utilized to maintain macro- and microsegmentation. As a part of that, the question of policy enforcement should be considered within the DMVPN environment. Will the DMVPN environment simply propagate the segmentation information, will only the headends enforce policy for traffic egressing to the hub locations, or will all DMVPN devices be responsible for policy enforcement?

Based on the answer to this question, configuring policy enforcement on a specific DMVPN device may be required. For instance, if DMVPN is providing enforcement only at the headend locations, then only the hubs will require specific enforcement configuration while the spokes will need only the propagation pieces. This is most likely the most common scenario, as traffic passing between SDA locations will have policy enforcement at the fabric egress point, such as the SDA fabric edge node. There will be migration scenarios where DMVPN is connecting SDA fabric sites with non-SDA sites, including the hub locations.

SDA and DMVPN Best Practices

For migrations to SDA with DMVPN, it would be expected that the DMVPN topology for a single network already exists in the enterprise, and the remote locations are transitioning to SDA. For this scenario, it is recommended to stand up the additional tunnel or tunnels in DMVPN at the hub and required spoke sites first. Routing and connectivity should be validated first prior to the SDA conversion. Here, standardizing tunnel numbers with SDA virtual networks is a good practice to facilitate easier management, operations, and troubleshooting.

Additionally, utilize the existing DMVPN tunnel for the SDA underlay topology. This will ensure that during migration, when devices are being configured in SDA, there will always be connectivity from the SDA underlay management interface with the Catalyst Center in the event that routing or connectivity issues occur in the newly deployed DMVPN tunnels. The overlay traffic will move into the new tunnel topology, leaving just the SDA underlay traffic in the original tunnel system.

DMVPN and ACI

DMVPN with ACI is similar to SD-WAN with ACI with the exception that there is no integration of the APIC with another controller. Just as with the SD-WAN-ACI integration, it is normally expected that macrosegmentation will be maintained between DMVPN and ACI, whereas microsegmentation propagation may not be a requirement. However, just as with the former, microsegmentation propagation may be configured, but again, with the limitation that the EPG mapping is limited to a single ACI context.

As seen in the DMVPN-SDA discussion, multiple DMVPN tunnels will be required on each DMVPN speaker with one tunnel per VRF required for macrosegmentation. It is recommended to use a fVRF for the tunnel source to reduce the complexity associated with ensuring tunnel route recursion does not occur.

Additionally, as discussed in the “SD-WAN and ACI” section earlier, the use of an aggregation layer between the ACI, DMVPN, and legacy environments will facilitate smoother migrations to either ACI, multi-tenant DMVPN, or both.

ACI and DMVPN Segmentation

It is possible with DMVPN to extend both the macrosegmentation and the microsegmentation created in ACI across the WAN topology; however, both are more complicated to configure and manage via the CLI as compared to the Cisco SD-WAN solution.

For macrosegmentation, the DMVPN environment may be made VRF aware by creating multiple DMVPN tunnels on each router for each required VRF. For each tunnel system, a unique NHRP ID and tunnel key should be utilized to ensure the various tunnels remain isolated. The use of CLI templates will greatly reduce the management overhead.

Extending microsegmentation from the ACI environment to the DMVPN environment is much more complicated and requires more manual configuration. In the preceding chapter discussing SDA, the SDA integration with ACI was shown using ISE pxGrid. Via ISE, the ACI EPG information is translated to the Cisco TrustSec SGT information. Using SDA, the Cisco DNA Center facilitates the management of the ISE deployment. With DMVPN, the Cisco TrustSec SGTs would have to be manually administered on ISE and connected to the appropriate ACI EPGs. Additionally, all of the Cisco TrustSec configuration must be manually configured. This would include adding SGT propagation on the DMVPN tunnels themselves, as shown previously, and including SXP peerings to the ISE for each VRF managed.

ACI and DMVPN Best Practices

It is best practice to standardize the DMVPN tunnel numbering for each ACI context with the VLAN used to interconnect the two at the hub locations. If an ACI tenant does not need to be extended to a specific remote site, do not configure that tenant’s tunnel at that site. This will limit the number of IPsec SAs required on each router, as well as the number of NHRP registrations.

If macrosegmentation is to be extended to the remote locations, then the macrosegmentation is inherently maintained via the various VRF tunnels. Microsegmentation propagation must be configured, if desired, using inline tagging on the tunnel configuration. For the remote endpoints in DMVPN to support Cisco TrustSec, they will need to support SXP communication with ISE. The scaling of the latter should be carefully considered because one pair of ISE nodes supports only 200 SXP peerings, which are counted as per VRF per device. Therefore, 20 devices with 10 VRFs each can reach the maximum limitation easily. ISE can scale to support 800 total SXP peerings with eight ISE nodes; however, a better scaling mechanism would be to utilize an SXP reflector—a dedicated router capable of handling a higher quantity of SXP peerings. As shown in Figure 3-12, the user creates an IP-to-SGT mapping on the ISE Policy Administration Node (PAN). This also could have been created dynamically by ISE. The IP-SGT mapping is forwarded to the policy node with the SXP support enabled. The SXP Policy Service Node (PSN)

will forward the update to the configured peer SXP Reflectors. The SXP Reflector could be an ASR1002HX or a Catalyst 8500 to facilitate the scaling. On the SXP Reflector side, the SXP peerings are per-VRF; therefore, the SXP Reflector updates only the DMVPN remote endpoints that have the specific VRF. Updates are not sent to other locations. When the SXP Reflector system is used, the load is reduced on the ISE environment while increasing the scale limitations. This does increase the count of devices to manage; however, these additional routers may be considered as server functionality similar to BGP route reflectors because they do not need to participate in the data path.

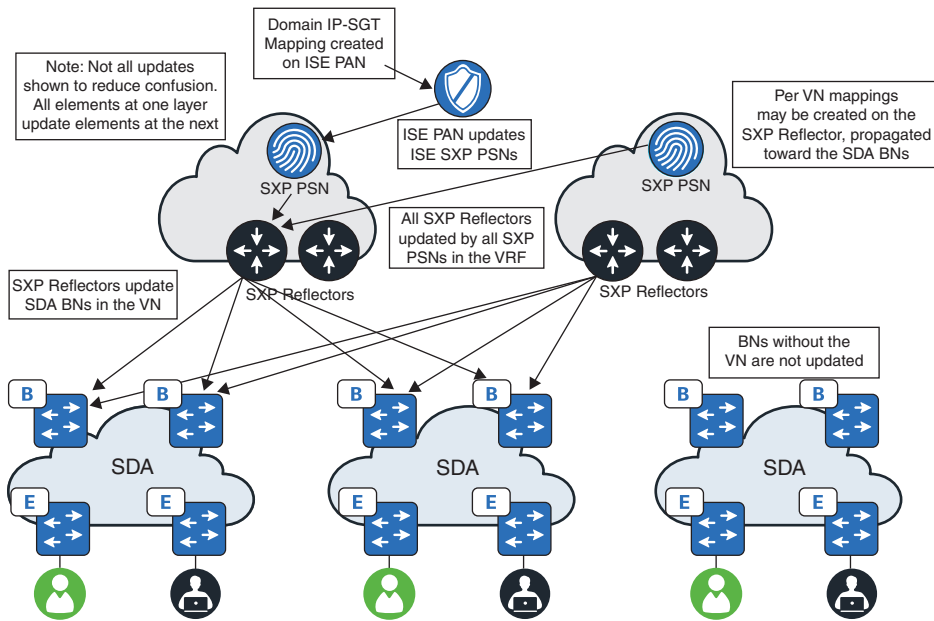


Figure 3-12 *SXP Reflector System*

Summary

Both Cisco SD-WAN and DMVPN solutions integrate well with the other domains, allowing the enterprise to extend the business intent and segmentation across the WAN environment between domains. While both solutions provide macrosegmentation via VRFs and microsegmentation by propagating the SGT value from one side of the WAN to the other, the management and configuration of the segmentation are quite different. For SD-WAN, the Catalyst SD-WAN Manager facilitates management of the SD-WAN Edge routers, whereas DMVPN requires manual configuration or use of another automation tool to manage the configurations. Additionally, DMVPN requires a unique tunnel system for each macrosegmented VRF, whereas the SD-WAN solution uses a single tunnel system with the ability to create logical topologies per VRF. In both solutions, having a single standard, such as VLAN-to-VRF mapping, used at all of the remote locations improves management and operational efficiencies.

Index

A

AAR (Application Aware Policies),
retail use case, 350–353

access

Internet access, retail use case,
353–355

NAC, 41, 42–43, 304

perimeter access, 41, 43–44

SDA, 22

ACI, best practices, 52–55

ACI, integrations, 31–34, 52–55

BN, 23–25

CNF integrations, 37–38

components of, 23–24

*Data Center Fabric
integrations, 31*

DMVPN, best practices, 64

DMVPN, configuring, 62–63

*DMVPN, policy
enforcement, 64*

FE nodes, 23–24

MPLS integrations, 34–37

MSRB, 40–41

NAC, 41, 42–43

peer devices, 23

perimeter access, 41, 43–44

SD-WAN, best practices, 51–52

SD-WAN, inline tagging, 50–51

SD-WAN, integrations, 30–31

SD-WAN, one-box SDA, 46

SD-WAN, segmentation, 49–51

*SD-WAN, two-box SDA
(nonintegrated solutions),
48–49*

security stacks, 41

*VNID, SDA and SD-WAN
segmentation, 49–51*

access layer architectures,
manufacturing use case, 290–293

access policies, ACI policy models,
79–80

ACI (Application Centric
Infrastructure)

ACI-Spine proxies, 84–86

AEP, 79–80

APIC, 74

application profiles, 78

automation, 282–284

- BD, 78
- Campus Fabric integrations, 31–34
- cloud (multi) configurations
 - cloud to External Enterprise Sites*, 154–161
 - cloud to Internet connections (direct connectivity)*, 150–151
 - cloud to Internet/External Destinations Using On-Prem L3Out Connectivity*, 151–153
 - guidelines/limitations*, 161–162
 - Inter-Tenant/Inter-VRF*, 148–150
 - Intra-Tenant/Intra-VRF*, 146–147
- cloud services (PaaS), 378
- declarative control, 70–71
- DMVPN, 64–65
 - best practices*, 65–66
 - segmentation*, 65
 - SXP Reflectors*, 65–66
- domains (different areas), 79–80
- endpoint groups, 78
- endpoint learning, 83–84
- enterprise MPLS, 178
 - super core MPLS*, 180–181
 - VRF*, 178
- EPG
 - access policies*, 79–80
 - tenant policies*, 78–79
- floods, 85
- forwarding, VXLAN overlays, 81–82
- FTAG, 85
- imperative control, 70–71
- interface policies, 80
- interface profiles, 80
- interface selectors, 80
- L4–7 policies, 80–81
- leaf switches, 83–85
- leaf-spine fabric design, 72–75
- logical design, 401–404
- MO, 77
- MPLS-SR integration, 388–404
- MPLS/SR-MPLS, pairwise integration, 111–132
- NDO, 75–76
- OpFlex control protocol, 70–72
- overview, 67–70
- pairwise integration
 - MPLS/SR-MPLS*, 111–132
 - SDA-ACI*, 101–110
 - SD-WAN*, 101–110
- policy models, 77
 - access policies*, 79–80
 - tenant policies*, 77–79
- public cloud integration, 132–135
- public sector use case
 - logical design*, 401–404
 - MPLS-SR integration*, 388–400
 - VMM integration*, 405–413
- retail use case, 340–343
- routing
 - COOP*, 83–85
 - VXLAN overlays*, 81–82
- SDA
 - best practices*, 52–55
 - integrations*, 31–34, 52–55
- SDA-ACI
 - group/identity automated mapping federation*, 89–93
 - integration guidelines/limitations*, 101
 - pairwise integration*, 88
 - VN-to-VRF automated mapping federation*, 93–95
 - VN-to-VRF manual mapping federation*, 96

- SD-WAN
 - guidelines/limitations*, 111
 - pairwise integration*, 101–110
 - segmentation*, 56
- security, 270–271
- switch profiles, 80
- switch selectors, 80
- traffic flows
 - ALB*, 144–146
 - AWS TGW*, 138–143
 - IPsec tunnels with VNG*, 143–144
 - IPsec tunnels with VGW*, 138–139
 - route programming in user VPC*, 142–143
 - VN peering*, 144–145
- underlay/overlay intersite connectivity, 135–138
- VMM integration, 405–413
- VRF, tenant policies, 78
- ACI-to-ACI Intersite (multi-site DC)**
 - connectivity Using SD-WAN, 106–110
- advertising routes to remote sites, retail use case, 348–349
- AEP (Attachable Access Entity Profiles)**, ACI, 79–80
- ALB (Azure Load Balancers)**, traffic flows, 144–146
- APIC (Application Policy Infrastructure Controllers)**
 - Catalyst SD-WAN Manager integration, 55–56
 - leaf-spine fabric design, 74
- application profiles**, ACI, 78
- applications**
 - AAR, retail use case, 350–353
 - cloud-based, 249–250
- Applications Stretched Across Sites**, **ACI multi-cloud configurations**
 - Inter-Tenant/Inter-VRF, 148–150
 - Intra-Tenant/Intra-VRF, 146–147
- architectures**
 - Kubernetes architectures, 410–413
 - manufacturing use case
 - access layer architectures*, 290–293
 - core and distribution layer architectures*, 293–295
 - wireless architectures*, 295–296
 - public sector use case, overall design architecture, 382–387
 - retail use case, overall design architecture, 328–330
 - transportation use case, overall design architecture, 419–421
- ASR 1000 series routers**, SDA, 23
- authentication**
 - dot1x authentication, 42–43
 - MAB, 42–43
- automated mapping**
 - group/identity automated mapping federation, SDA-ACI, 89–93
 - VN-to-VRF automated mapping federation, SDA-ACI, 93–95
- automation**, 19
 - ACI, 282–284
 - CI/CD pipelines, 273–275
 - DMVPN, 281–282
 - manufacturing use case, 301
 - multiple domains, 284
 - overview, 273
 - public sector use case, 413–416
 - SDA, 276–278
 - SD-WAN, 278–281
- availability versus downtime, redundancy**, 3–4

AWS (Amazon Web Service), TGW traffic flows, 138–143

Azure service deployments, retail use case, 375–377

B

backstage sites, retail use case, 357–358

back-to-back VRF, enterprise MPLS, ACI, 178

BD (Bridge Domains), ACI, 78

best practices

DMVPN and ACI, 65–66

DMVPN and SDA, 64

SDA

ACI integrations, 56

SD-WAN, 51–52

BGP-EVPN, SDA and ACI integrations, 32–34

BMS (Building Management Systems)

BU, 15

network segmentation, 18

traffic isolation, 17

trust boundaries, 13

BN (Border Nodes)

SDA, 23–25

virtual network handoffs, 23–25

borders

CNF, dedicated borders, 196–199

MSRB, CNF, 196–199

boundaries

IT boundary, 16–17

trust boundaries, 13

branch networks, retail use case, 355–359

BU (Business Units), 15–16

business requirements, manufacturing use case, 287–288

C

Campus Fabric, 22, 25

ACI integrations, 31–34

CNF, 193–194

dedicated borders, 196–199

integrations, 37–38

MSRB, 196–199

VRF-Lite, 194–195

MPLS integrations, 34–37

NAC, 41, 42–43

perimeter access, 41, 43–44

SD-WAN, 25–26

default end-to-end

macrosegmentation, 26–27

integrated/non-integrated

SD-WAN solutions, 27–30

one-box solutions, 27, 29–30

route architectures, 28–30

SDA integrations, 30–31

two-box solutions, 27, 28

security stacks, 41

campus networks

Campus Fabric, 22, 25

SD-WAN, 25–26

SD-WAN, default end-to-end

macrosegmentation, 26–27

SD-WAN, integrated/non-

integrated SD-WAN solutions, 27–30

SD-WAN, one-box solutions, 27, 29–30

SD-WAN, route architectures, 28–30

SD-WAN, SD-Access integrations, 30–31

SD-WAN, two-box solutions, 27, 28

cloud computing

private clouds, 233–235*public clouds*, 235–237

financial use case

data center services, SDA,
318–320*data center services*, SD-WAN,
321*extending segmentation*,
315–317*macrosegmentation*, 310–312*microsegmentation*, 312–315*modular services*, 308–310*SDA virtual networks*, 312*SD-WAN centralized policies*,
322–326

retail use case, 355–359

SDA, 22

ACI integrations, 31–34*BN*, 23–25*CNF integrations*, 37–38*components of*, 23–24*Data Center Fabric
integrations*, 31*FE nodes*, 23–24*MPLS integrations*, 34–37*MSRB*, 40–41*NAC*, 41, 42–43*peer devices*, 23*perimeter access*, 41, 43–44*SD-WAN integrations*, 30–31*security stacks*, 41

transportation use case, 425–426

macrosegmentation, 429–432*microsegmentation*, 432–433*transportation-specific entities*,
426–429

Catalyst 9K Layer 3 switches, SDA, 23

Catalyst 8000 series routers, SDA, 23

Catalyst Center, SD-WAN
integrations, 48Catalyst SD-WAN Manager, 48,
55–56

CI/CD pipelines, 273–275

cloud computing

ACI multi-cloud configurations

*cloud to External Enterprise
Sites*, 154–161*cloud to Internet connections
(direct connectivity)*, 150–151*cloud to Internet/External
Destinations Using On-Prem
L3Out Connectivity*, 151–153*guidelines/limitations*, 161–162*Intra-Tenant/Intra-VRF*,
146–147, 148–150

applications, 249–250

campus networks

private clouds, 233–235*public clouds*, 235–237

Cloud OnRamp, SD-WAN, 238–243

CNF, 220–224

private clouds, 243–245*public clouds*, 245

CoR, SD-WAN, 59–61

data/file storage, 249

email, 248–249

hybrid clouds, 232–233, 371–375

malware, 249–250

MPLS, 246–247

overview, 229–231

private clouds, 231–232

campus networks, 233–235*CNF*, 243–245*MPLS*, 246–247

- public clouds, 232
 - ACI integration*, 132–135
 - campus networks*, 235–237
 - CNF, 245
 - SD-WAN, 57–58, 267
 - SD-WAN Edge, 57–58
 - security, 271–272
 - security stacks, 247–251
 - viruses, 249–250
 - visibility, 250
 - cloud edge, retail use case, 337–338
 - cloud services (PaaS), ACI
 - integration, 378
 - cloud to Internet/External
 - Destinations Using On-Prem
 - L3Out Connectivity, 151–153
 - CNF (Carrier Neutral Facilities), 37–38, 199–200
 - Campus Fabric, 193–194
 - dedicated borders*, 196–199
 - integrations*, 37–38
 - MSRB, 196–199
 - VRF-Lite, 194–195
 - cloud computing, 220–224
 - private clouds*, 243–245
 - public clouds*, 245
 - data centers, 214–217
 - dark fiber connections*, 217–219
 - dedicated borders*, 196–199
 - WAN connections, 220
 - enterprise MPLS, 181
 - inter-VRF routing/merging services*, 182–184
 - NNI connections*, 184–187
 - offloading selective VRF*, 181–182
 - manufacturing use case, 298–301
 - overview, 191–193
 - retail use case, 330–332
 - SDA, 193–194
 - dedicated borders*, 196–199
 - MSRB, 196–199
 - VRF-Lite, 194–195
 - SD-WAN, 200–202
 - DMVPN, 202–204
 - E-MPLS, 209–212
 - P-MPLS, 206–209
 - security stacks, 224–227
 - VRF-Lite, 194–195
 - CNF Backbone, retail use case, 338–340
 - control plane, 22
 - COOP (Council of Oracles Protocol)
 - ACI routing/forwarding, 83–85
 - FTAG, 85
 - leaf switches, 83–85
 - CoR (Cloud OnRamp), SD-WAN, 59–61
 - core and distribution layer
 - architectures, manufacturing use case, 293–295
- ## D
-
- dark fiber, CNF/data center
 - connections, 217–219
 - Data Center Fabric, SDA
 - integrations, 31
 - data centers
 - CNF, 214–217
 - dark fiber connections*, 217–219
 - dedicated borders*, 196–199
 - WAN connections, 220
 - redundancy, 8
 - retail use case, 330–332

- services
 - financial use case*, 318–321
 - transportation use case*, 433–437
- transportation use case, 422–425, 431–432
- data plane, 22
- data/file storage, cloud computing, 249
- DC and Branch Using SD-WAN, 104–108
- declarative control
 - ACI, 70–71
 - SDN, 68
- dedicated borders, CNF, 196–199
- default end-to-end
 - macrosegmentation, 26–27
- deployed solutions, manufacturing use case, 290–291
- device redundancy, 5–7
- different areas (domains), 13, 79–80
- direct cloud to Internet connections, 150–151
- DMVPN (Dynamic Multi-Point VPN), 34, 45, 61–62
 - ACI, 64–65
 - best practices*, 65–66
 - segmentation*, 65
 - SXP Reflectors*, 65–66
 - automation, 281–282
 - CNF SD-WAN, 202–204
 - SDA
 - best practices*, 64
 - configuring*, 62–63
 - policy enforcement*, 64
 - security, 267–270
 - segmentation, 18–19
- DMZ (Demilitarized Zones)
 - manufacturing use case, 302–305

- MPLS-SR, ACI integration, 395–400
- SD-WAN security, 262–267
- DN (Distinguished Names), MO, 77
- domains (different areas), 13
 - ACI, 79–80
 - fault domains, 13–14
- dot1x authentication, 42–43
- downtime versus availability, redundancy, 3–4
- Dynamic Tunnel Control Policy, transportation use case, 437–440

E

- east-west SGT enforcements between endpoints
 - in different sites, 367–369
 - in the same site, 362–367
- email, cloud computing, 248–249
- E-MPLS (Enterprise MPLS), CNF SD-WAN, 209–212
- endpoint databases, 84
- endpoint groups, ACI, 78
- endpoints
 - identity, 84
 - learning, ACI, 83–84
 - locations, 84
- end-to-end macrosegmentation, default, 26–27
- end-to-end segmentation, SDA-Transit, 38–40
- enterprise MPLS
 - ACI, 178
 - super core MPLS*, 180–181
 - VRF*, 178
 - CNF, 181
 - inter-VRF routing/merging services*, 182–184

- NNI connections*, 184–187
 - offloading selective VRF*, 181–182
- deploying, 34
- overview, 163–166
- SDA, 166–167
 - default integration*, 167–168
 - integrations*, 34
 - SXP, SGT distributions*, 169–170
- SD-WAN, 170–171
 - super core MPLS*, 172–178
 - WAN transport*, 171–172
- security stacks, 187–188
- super core MPLS
 - ACI*, 180–181
 - SD-WAN*, 172–178
- VPN route-leaking, 183–184
- WAN transport, 171–172
- entry/exit points, 12–13
- EPG (Endpoint Group Tags), 18
 - ACI
 - access policies*, 79–80
 - tenant policies*, 78–79
 - SDA and ACI integrations, 31–34
- External Enterprise Sites, cloud (multi) configurations, 154–161
- external firewalls, SDA fabric security, 257–258
 - perimeter access*, 41, 43–44
 - SD-WAN*, 25–31
 - security stacks*, 41
- SDA, 257–261
- fault domains, 13–14
- FE (Fabric Edge) nodes, 23–24
- file/data storage, cloud computing, 249
- financial use case
 - campus networks
 - data center services, SDA*, 318–320
 - data center services, SD-WAN*, 321
 - extending segmentation*, 315–317
 - macrosegmentation*, 310–312
 - microsegmentation*, 312–315
 - modular services*, 308–310
 - SDA virtual networks*, 312
 - SD-WAN centralized policies*, 322–326
 - overview, 307–308
- firewalls
 - FTD firewall, SDA, 23
 - retail use case
 - trusted firewalls*, 335–337
 - untrusted firewalls*, 333–334
 - SDA fabric security, 257–258
- floods, ACI, 85
- forwarding
 - ACI
 - COOP*, 83–85
 - VXLAN overlays*, 81–82
 - VRF, 17–18, 22, 23
 - ACI tenant policies*, 78
 - back-to-back VRF, enterprise MPLS ACI*, 178
 - enterprise MPLS, ACI*, 178

F

fabrics

- Campus Fabric, 22, 25
 - ACI integrations*, 31–34
 - CNF*, 37–38, 193–199
 - MPLS integrations*, 34–37
 - NAC*, 41, 42–43

FTD, SDA, 23

*inter-VRF routing/merging
services, enterprise MPLS,
182–184*

*MPLS-SR, ACI integration,
390–392, 393–395*

*offloading selective VRF,
enterprise MPLS, 181–182*

*VN-to-VRF automated
mapping federation,
SDA-ACI, 93–95*

*VN-to-VRF manual mapping
federation, SDA-ACI, 96*

VRF-Lite, 23, 194–195

FTAG (Forwarding TAGs), ACI-COOP
routing/forwarding, 85

FTD (FirePower Threat Defense)
firewall, SDA, 23

fusion routers, 23

G

geo-redundancy, 8

group/identity automated mapping
federation, SDA-ACI, 89–93

H

Headend Control Policy, financial use
case, 322–324

headends (SD-WAN), retail use case,
334–335

hybrid clouds, 232–233, 371–375

I

IBN (Intent-Based Networking), 18

identity, group/identity automated
mapping federation (SDA-ACI),
89–93

imperative control

ACI, 70–71

SDN, 68

inline tagging

SDA and SD-WAN segmentation,
50–51

SGT, 258–259

**integrated solutions. *See*
one-box SDA**

**integrated/non-integrated SD-WAN
solutions, 27–30**

interface policies, ACI, 80

interface profiles, ACI, 80

interface selectors, ACI, 80

Internet Edge, retail use case, 333

**Inter-Tenant/Inter-VRF ACI multi-
cloud configurations, 148–150**

**Intra-Tenant/Intra-VRF ACI multi-
cloud configurations, 146–147**

**IP A addresses, ACI-COOP routing/
forwarding, 83**

IPsec tunnels

with VNG, 143–144

with VGW, 138–139

**ISE (Identity Services Engine), SDA
and ACI integrations, 31–34**

ISE TrustSec matrix, 260–261

isolating traffic, 17–18

IT boundary, 16–17

**IT DMZ, manufacturing use case,
302–304**

**ITaaS (Information Technology as a
Business), 14–15, 163–165**

BU, 15–16

IT boundary, 16–17

**iVXLAN (Intelligent VXLAN), ACI
routing/forwarding, 82**

J - K - L

Kubernetes architectures, 410–413

L4–7 policies, ACI, 80–81

LAN (Local-Area Networks)

segmentation, 18–19

VLAN, retail use case

backstage sites, 357–358

retail sites, 359

VXLAN

ACI routing/forwarding, 81–82

*iVXLAN, ACI routing/
forwarding, 82*

*SDA and ACI integrations,
31–34*

Layer 3 switches, SDA, 23

leaf switches, ACI-COOP routing/
forwarding, 83–85

leaf-spine fabric design, 72, 74–75

ACI-Spine proxies, 84–86

APIC, 74

link redundancy, 7

load balancing, ALB traffic flows,
144–146

LPM tables, ACI-COOP routing/
forwarding, 83

M

MAB (MAC Authentication Bypass),
42–43

MAC A addresses, ACI-COOP
routing/forwarding, 83

MAC addresses, MAB, 42–43

macrosegmentation

default end-to-end

macrosegmentation, 26–27

financial use case, 310–312

transportation use case, 429–432

VN-to-VRF automated mapping
federation, SDA-ACI, 93–95

VN-to-VRF manual mapping
federation, SDA-ACI, 96

malware, cloud computing, 249–250

management plane, 22

manufacturing use case

automation, 301

CNF, 298–301

deployed solutions, 290–291

DMZ, 302–305

IT DMZ, 302–304

microsegmentation, 304

NAC, 304

OT DMZ, 303–305

overview, 285

perimeter security, 302–303

requirements

business requirements, 287–288

summary of, 286–287

technical requirements, 288–289

SDA, 290–292

*access layer architectures,
290–293*

*core and distribution layer
architectures, 293–295*

wireless architectures, 295–296

SD-WAN, 297–298

security, 301–304

solution mapping

business requirements, 287–288

technical requirements, 288–289

mapping

databases, 84

group/identity automated mapping
federation, SDA-ACI, 89–93

- SGT
 - north-south flows*, 369
 - SIG traffic*, 369
- solution mapping
 - manufacturing use case*,
business requirements,
287–288
 - manufacturing use case*,
technical requirements,
288–289
- virtual networks, transportation use case, 431–432
- VN-to-VRF automated mapping federation, SDA-ACI, 93–95
- VN-to-VRF manual mapping federation, SDA-ACI, 96
- Meraki, retail use case**, 359–362
- MGMT, retail use case**, 338
- microsegmentation**
 - group/identity automated mapping federation, SDA-ACI, 89–93
 - manufacturing use case, 304
 - transportation use case, 432–433
- MO (Managed Objects), ACI**, 77
- modular services**
 - financial use case, 308–310
 - microsegmentation, 312–315
- MPLS (Multi-Protocol Label Switching)**
 - ACI, 178
 - pairwise integration*, 111–132
 - super core MPLS*, 180–181
 - VRF*, 178
 - Campus Fabric integrations, 34–37
 - cloud computing, 246–247
 - CNF, 181
 - inter-VRF routing/merging services*, 182–184
 - NNI connections*, 184–187
 - offloading selective VRF*,
181–182
 - E-MPLS, CNF SD-WAN, 209–212
 - overview, 163–166
 - P-MPLS, CNF SD-WAN, 206–209
 - SDA, 166–167
 - default integration*, 167–168
 - integrations*, 34–37
 - SXP, SGT distributions*,
169–170
 - SDA-Transit over MPLS, 36–37
 - SD-WAN, 56–57, 170–171
 - super core MPLS*, 172–178
 - WAN transport*, 171–172
 - security stacks, 187–188
 - super core MPLS
 - ACI*, 180–181
 - SD-WAN*, 172–178
 - VPN route-leaking, 183–184
 - WAN transport, 171–172
- MPLS Edge, retail use case**, 333
- MPLS-SR (MPLS-Segment Routing)**
 - ACI integration, 111–132,
388–404
 - DMZ, 395–400
 - pairwise integration, 111–132
 - public sector use case, 382–387
 - VPN, 386–387
 - MSO (Multi-Site Orchestrator).
See NDO
- MSRB (Multi-Site Remote Borders)**,
40–41, 196–199
- multi-cloud ACI configurations**
 - cloud to Internet connections (direct connectivity), 150–151

- cloud to Internet/External Destinations Using On-Prem L3Out Connectivity, 151–153
- Intra-Tenant/Intra-VRF, 146–147, 148–150
- multi-connectivity, 9–10**
- multi-domain networks**
 - automation, 19
 - characteristics of, 11–14
 - components of, 3–11
 - defined, 2
 - domains (different areas), 13–14
 - entry/exit points, 12–13
 - fault domains, 13–14
 - need for, 14
 - orchestration, 19
 - redundancy, 3
 - availability versus downtime, 3–4*
 - device redundancy, 5–7*
 - geo-redundancy, 8*
 - link redundancy, 7*
 - multi-connectivity, 9–10*
 - multiple technologies, 11*
 - power redundancy, 5*
 - resiliency versus redundancy, 9*
 - shared services, 10–11*
 - site/data center redundancy, 8*
 - types of, 4–8*
 - WAN, 9–10*
 - trust boundaries, 13
- multiple domains, automation, 284**
- multiple technologies, 11**
- multi-site DC (ACI-to-ACI Intersite) connectivity Using SD-WAN, 106–110**
- multi-tenancy, 17–18**

N

- NAC (Network Access Control), 41, 42–43, 304**
- NDO (Nexus Dashboard Orchestrator), 75–76**
 - VN-to-VRF automated mapping federation, SDA-ACI, 93–95
 - VN-to-VRF manual mapping federation, SDA-ACI, 96
- Nexus 9K Layer 3 switches, SDA, 23**
- NNI (Network-to-Network Interfaces), enterprise MPLS, 184–187**
- non-integrated/integrated SD-WAN solutions, 27–30**
- north-south flows, SGT mapping, 369**

O

- one-box SDA (integrated solutions), SD-WAN, 46**
- one-box solutions, 27, 29–30**
- On-Prem L3Out Connectivity, Cloud to Internet/External Destinations Using, 151–153**
- OpFlex control protocol, 70–72**
- orchestration, 19, 413–416**
- OT DMZ, manufacturing use case, 303–305**
- overlays**
 - intersite connectivity, 135–138
 - SD-WAN, security, 262–267

P

- PaaS (cloud services), ACI integrations, 378**
- pairwise integration**
 - MPLS/SR-MPLS ACI, 111–132

- SDA-ACI, 88
- SD-WAN ACI, 101–108
- PBR (Policy-Based Routing), L4–7
 - policies, 80–81
- peer devices, 23
- peering, VN, 144–145
- perimeters
 - access, 41, 43–44
 - security, manufacturing use case, 302–303
- P-MPLS (Provider-MPLS), CNF
 - SD-WAN, 206–209
- policies
 - AAR, retail use case, 350–353
 - enforcement, DMVPN and SDA, 64
 - financial use case
 - Headend Control Policy*, 324–326
 - Site ID-Based Country Control Policy*, 322–324
 - retail use case
 - Meraki*, 359–362
 - TrustSec*, 359–362
 - SD-WAN centralized policies
 - financial use case*, 322–326
 - transportation use case*, 437–440
 - security, 254
 - transportation use case, Dynamic Tunnel Control Policy, 437–440
 - zone-pair policies, SD-WAN security, 264–267
- policy models, ACI, 77
 - access policies, 79–80
 - tenant policies, 77–79
- policy plane, 22
- power redundancy, 5
- private clouds, 231–232

- campus networks, 233–235
- CNF, 243–245
- MPLS, 246–247
- public clouds, 232
 - ACI integration, 132–135
 - campus networks, 235–237
 - CNF, 245
- public sector use case
 - ACI
 - logical design*, 401–404
 - MPLS-SR integration*, 388–400
 - VMM integration*, 405–413
 - automation, 413–416
 - MPLS-SR, 382–386
 - ACI integration*, 388–400
 - DMZ*, 395–400
 - VPN*, 386–387
 - orchestration, 413–416
 - overall design architecture, 382–387
 - overview, 381–382

Q

- QoS (Quality of Service)
 - classifications and markings, retail use case, 352–353

R

- receiving routes from remote sites, retail use case, 348–351
- redundancy, 3
 - availability versus downtime, 3–4
 - device redundancy, 5–7
 - geo-redundancy, 8
 - link redundancy, 7
 - multi-connectivity, 9–10

- multiple technologies, 11
- power redundancy, 5
- resiliency versus redundancy, 9
- shared services, 10–11
- site/data center redundancy, 8
- types of, 4–8
- WAN, 9–10
- remote sites, retail use case**
 - advertising routes to remote sites, 348–349
 - receiving routes from remote sites, 348–351
- resiliency versus redundancy, 9**
- retail use case**
 - AAR, 350–353
 - ACI, 340–343
 - advertising routes to remote sites, 348–349
 - Azure service deployments, 375–377
 - backstage sites, 357–358
 - branch networks, 355–359
 - campus networks, 355–359
 - cloud edge, 337–338
 - cloud services (PaaS), ACI integration, 378
 - CNF, 330–332
 - CNF Backbone, 338–340
 - connectivity, 375–377
 - data centers, 330–332
 - hybrid cloud integration, 371–375
 - Internet Edge, 333
 - Meraki, 359–362
 - MGMT, 338
 - MPLS Edge, 333
 - overall design architecture, 328–330
 - overview, 327–328
 - QoS classifications and markings, 352–353
 - receiving routes from remote sites, 348–351
 - retail sites, 359
 - SD-WAN, 343–347
 - headends*, 334–335
 - VPN, 350
 - secure Internet access, 353–355
 - SGT, 359–369
 - SIG, 353–355
 - TLOC tunnel groups/color, 347
 - trusted firewalls, 335–337
 - TrustSec, 359–362
 - untrusted firewalls, 333–334
- route programming, VPC, 142–143**
- route-leaking, VPN, 183–184**
- routing**
 - ACI
 - COOP, 83–85
 - VXLAN overlays, 81–82
 - ASR 1000 series routers, SDA, 23
 - Catalyst 8000 series routers, SDA, 23
 - fusion routers, 23
 - MPLS-SR
 - ACI integration, 388–404
 - DMZ, 395–400
 - public sector use case*, 382–387
 - VPN, 386–387
 - PBR, L4–7 policies, 80–81
 - retail use case
 - advertising routes to remote sites*, 348–349
 - receiving routes from remote sites*, 348–351
 - SD-WAN
 - route architectures*, 28–30
 - routers*, SDA, 23
 - VRF, 17–18, 22, 23

- ACI tenant policies, 78*
- back-to-back VRF, enterprise MPLS ACI, 178*
- enterprise MPLS, ACI, 178*
- FTD, SDA, 23*
- inter-VRF routing/merging services, enterprise MPLS, 182–184*
- MPLS-SR, ACI integration, 390–392, 393–395*
- offloading selective VRF, enterprise MPLS, 181–182*
- VN-to-VRF automated mapping federation, SDA-ACI, 93–95*
- VN-to-VRF manual mapping federation, SDA-ACI, 96*
- VRF-Lite, 23, 194–195*

S

SASE (Secure Access Service Edge), SIG, 58–59

SDA (SD-Access), 22

ACI

- best practices, 52–55*
- group/identity automated mapping federation, 89–93*
- integration guidelines/limitations, 101*
- integrations, 31–34, 52–55*
- pairwise integration, 88*
- VN-to-VRF automated mapping federation, 93–95*
- VN-to-VRF manual mapping federation, 96*
- automation, 276–278
- BN, 23–25
- CNF, 193–194

- dedicated borders, 196–199*
- integration, 37–38*
- MSRB, 196–199*
- VRF-Lite, 194–195*
- components of, 23–24
- Data Center Fabric integrations, 31
- DMVPN
 - best practices, 64*
 - configuring, 62–63*
 - policy enforcement, 64*
- enterprise MPLS, 166–167
 - default integration, 167–168*
 - SXP, SGT distributions, 169–170*
- FE nodes, 23–24
- financial use case
 - campus networks, data center services, 318–320*
 - virtual networks, 312*
- manufacturing use case, 290–292
 - access layer architectures, 290–293*
 - core and distribution layer architectures, 293–295*
 - wireless architectures, 295–296*
- MPLS integrations, 34–37
- MSRB, 40–41
- NAC, 41, 42–43
- peer devices, 23
- perimeter access, 41, 43–44
- SD-WAN, 46
 - best practices, 51–52*
 - inline tagging, 50–51*
 - integration, 30–31*
 - one-box SDA (integrated solutions), 46*
 - segmentation, 49–51*
 - two-box SDA (nonintegrated solutions), 48–49*

- security, 254–256
 - fabric*, 257–261
 - users*, 254–256
- security stacks, 41
- transportation use case
 - bus SGT*, 432–433
 - data center services*, 433–437
 - macrosegmentation*, 429–431
- virtual networks, 431
- VNID, SDA and SD-WAN
 - segmentation, 49–51
- YAML, 277–278
- SDA-Transit**
 - end-to-end segmentation, 38–40
 - MPLS integrations, 36–37
- SDN (Software-Defined Networking)**
 - controllers, 67–68
 - declarative control, 68
 - imperative control, 68
 - overview, 67–70
- SD-WAN (Software-Defined WAN), 45**
 - ACI
 - guidelines/limitations*, 111
 - pairwise integration*, 101–110
 - segmentation*, 56
 - automation, 278–281
 - Campus Fabric, 25–26, 28
 - default end-to-end macrosegmentation*, 26–27
 - integrated/non-integrated SD-WAN solutions*, 27–30
 - one-box solutions*, 27, 29–30
 - SDA integrations*, 30–31
 - SD-WAN route architectures*, 28–30
 - two-box solutions*, 27
 - Catalyst Center integrations, 48
 - Catalyst SD-WAN Manager, 48, 55–56
 - centralized policies, financial use case, 322–326
 - cloud computing, 57–58
 - Cloud OnRamp, 238–243
 - cloud to Internet/External Destinations Using On-Prem L3Out Connectivity, 157–161
 - CNF, 200–202
 - DMVPN*, 202–204
 - E-MPLS*, 209–212
 - P-MPLS*, 206–209
 - CoR, 59–61
 - DC and Branch Using SD-WAN, 104–108
 - enterprise MPLS, 170–171
 - super core MPLS*, 172–178
 - WAN transport*, 171–172
 - financial use case, campus networks
 - data center services*, 321
 - SD-WAN centralized policies*, 322–326
 - Headend Control Policy, 324–326
 - integrated/non-integrated SD-WAN solutions, Campus Fabric, 27–30
 - manufacturing use case, 297–298
 - MPLS, 56–57
 - multi-site DC (ACI-to-ACI Intersite) connectivity Using SD-WAN, 106–110
 - retail use case, 343–347
 - SD-WAN headends*, 334–335
 - VPN*, 350
 - SDA, 23, 46
 - best practices*, 51–52
 - inline tagging*, 50–51
 - one-box SDA (integrated solutions)*, 46

- segmentation*, 49–51
- two-box SDA (nonintegrated solutions)*, 48–49
- security
 - cloud computing*, 267
 - DMZ zones*, 262–267
 - overlays*, 262–267
 - underlays*, 262
 - zone-pair policies*, 264–267
- segmentation, 18–19
- SIG, 58–59
- Site ID-Based Country Control Policy, 322–324
- templates, 279–281
- transportation use case, SD-WAN
 - centralized policies, 437–440
- SD-WAN Edge**, 57–58
- security
 - ACI, 270–271
 - cloud computing, 249–250, 271–272
 - applications*, 249–250
 - data/file storage*, 249
 - email*, 248–249
 - malware*, 249–250
 - viruses*, 249–250
 - visibility*, 250
 - DMVPN, underlays, 267–270
 - firewalls
 - SDA fabric security*, 257–258
 - trusted firewalls, retail use case*, 335–337
 - untrusted firewalls, retail use case*, 333–334
 - Internet access, retail use case, 353–355
 - manufacturing use case, 301–304
 - Meraki, retail use case, 359–362
 - overview, 253
 - policies, 254
 - SDA, 254–256
 - fabric*, 257–261
 - users*, 254–256
 - SD-WAN
 - cloud computing*, 267
 - DMZ zones*, 262–267
 - overlays*, 262–267
 - underlays*, 262
 - zone-pair policies*, 264–267
 - SGT inline tagging, 258–259
 - SXP reflectors, 259–260
 - TrustSec
 - matrix*, 260–261
 - retail use case*, 359–362
 - zero trust security, 272
- security stacks**
 - cloud computing, 247–251
 - CNF, 224–227
 - dot1x authentication, 42–43
 - enterprise MPLS, 187–188
 - MAB, 42–43
 - NAC, 41, 42–43
 - perimeter access, 41, 43–44
 - SDA, 41
- segmentation**
 - default end-to-end
 - macrosegmentation, 26–27
 - DMVPN and ACI, 65
 - end-to-end segmentation, SDA-Transit, 38–40
 - financial use case, extending segmentation, 315–317
 - networks, 18–19
 - SDA, SD-WAN segmentation, 49–51
- service graphs**, L4–7 policies, 80–81

SGT (Security Group Tags), 18
 assignments, 361–362
 east-west enforcements between endpoints
 in different sites, 367–369
 in the same site, 362–367
 inline tagging, 258–259
 mapping
 north-south flows, 369
 SIG traffic, 369
 retail use case, 359–369
 SDA and ACI integrations, 31–34
 SXP and SGT distributions, 169–170
 transportation use case, 432–433
shared services, 10–11
SIG (Secure Internet Gateway)
 retail use case, 353–355
 SD-WAN, 58–59
 SGT mapping, 369
site/data center redundancy, 8
Site ID-Based Country
 Control Policy, financial use case, 322–324
solution mapping, manufacturing use case
 business requirements, 287–288
 technical requirements, 288–289
spines-leaf fabric design, 72, 74–75
 ACI-Spine proxies, 84–86
 APIC, 74
SR-MPLS (Segment Routing MPLS)
 ACI integration, 111–132, 388–404
 DMZ, 395–400
 pairwise integration, 111–132
 public sector use case, 382–387
 VPN, 386–387
storage (data/files), cloud computing, 249

super core MPLS
 ACI, 180–181
 SD-WAN, 172–178
switch profiles, ACI, 80
switch selectors, ACI, 80
switches
 leaf switches, ACI-COOP routing/forwarding, 83–85
 SDA
 Catalyst 9K Layer 3 switches, 23
 Layer 3 switches, 23
SXP (Security Group Tag Exchange Protocol), SGT distributions, 169–170
SXP reflectors, 65–66, 259–260

T

tags
 EPG, 18
 ACI, 78–80
 SDA and ACI integrations, 31–34
 FTAG, ACI-COOP routing/forwarding, 85
 inline tagging, SDA and SD-WAN segmentation, 50–51
 SGT, 18
 assignments, 361–362
 east-west enforcements between endpoints, 362–369
 inline tagging, 258–259
 mapping, 369
 retail use case, 359–369
 SDA and ACI integrations, 31–34
 SXP and SGT distributions, 169–170
 transportation use case, 432–433

- TCP (Transit Control Plane),
dedicated borders, 196–199
- technical requirements, manufacturing
use case, 288–289
- templates, SD-WAN, 279–281
- tenant policies, ACI policy models,
77–79
- TGW (Transit Gateways), traffic
flows, 138–143
- TLOC (Transport Locator) tunnel
groups/color, 347
- traffic flows
 - ALB, 144–146
 - AWS TGW, 138–143
 - IPsec tunnels
 - with VNG, 143–144*
 - with VGW, 138–139*
 - route programming in user VPC,
142–143
 - VN peering, 144–145
- traffic isolation, 17–18
- transportation use case
 - campus networks, 425–426
 - macrosegmentation, 429–432*
 - microsegmentation, 432–433*
 - transportation-specific entities,*
426–429
 - data centers, 422–425, 431–432
 - Dynamic Tunnel Control Policy,
437–440
 - macrosegmentation, 429–432
 - microsegmentation, 432–433
 - overall design architecture, 419–421
 - overview, 419
 - SDA
 - bus SGT, 432–433*
 - data center services, 433–437*
 - macrosegmentation, 429–431*

- SD-WAN centralized policies,
437–440
- virtual networks
 - mapping, 431–432*
 - SDA, 431*
- WAN, macrosegmentation, 431–432
- trust boundaries, 13
- trusted firewalls, retail use case,
335–337
- TrustSec
 - matrix, 260–261
 - retail use case, 359–362
- tunneling
 - Dynamic Tunnel Control Policy,
437–440
 - IPsec tunnels
 - with VGW, 138–139*
 - with VNG, 143–144*
 - TLOC tunnel groups/color, 347
- two-box SDA (nonintegrated
solutions), SD-WAN, 48–49
- two-box solutions, 27, 28

U

- underlays
 - intersite connectivity, 135–138
 - SD-WAN
 - security, 262*
 - zone-pair policies, 264–267*
- untrusted firewalls, retail use case,
333–334
- use cases
 - financial use case
 - campus networks, data center*
services, 318–321
 - campus networks, extending*
segmentation, 315–317

- campus networks, macrosegmentation*, 310–312
- campus networks, microsegmentation*, 312–315
- campus networks, modular services*, 308–310
- campus networks, SDA*, 318–320
- campus networks, SDA virtual networks*, 312
- campus networks, SD-WAN*, 321
- campus networks, SD-WAN centralized policies*, 322–326
- overview*, 307–308
- group/identity automated mapping federation, SDA-ACI, 89–93
- manufacturing use case, 285
 - automation*, 301
 - CNF, 298–301
 - deployed solutions*, 290–291
 - DMZ, 302–305
 - IT DMZ, 302–304
 - microsegmentation*, 304
 - NAC, 304
 - OT DMZ, 303–305
 - perimeter security*, 302–303
 - requirements, business requirements*, 287–288
 - requirements, summary of*, 286–287
 - requirements, technical requirements*, 288–289
 - SDA, 290–292
 - SDA, access layer architectures*, 290–293
 - SDA, core and distribution layer architectures*, 293–295
 - SDA, wireless architectures*, 295–296
 - SD-WAN, 297–298
 - security*, 301–304
 - solution mapping, business requirements*, 287–288
 - solution mapping, technical requirements*, 288–289
- public sector use case
 - ACI, *logical design*, 401–404
 - ACI, *MPLS-SR integration*, 388–400
 - ACI, *VMM integration*, 405–413
 - automation*, 413–416
 - MPLS-SR, 382–387
 - MPLS-SR, *ACI integration*, 388–400
 - MPLS-SR, *DMZ*, 395–400
 - MPLS-SR, *VPN*, 386–387
 - orchestration*, 413–416
 - overall design architecture*, 382–387
 - overview*, 381–382
- retail use case
 - AAR, 350–353
 - ACI, 340–343
 - advertising routes to remote sites*, 348–349
 - Azure service deployments*, 375–377
 - backstage sites*, 357–358
 - branch networks*, 355–359
 - campus networks*, 355–359
 - cloud edge*, 337–338
 - cloud services (PaaS), ACI integration*, 378
 - CNF, 330–332
 - CNF Backbone, 338–340
 - connectivity*, 375–377
 - data centers*, 330–332
 - hybrid cloud integration*, 371–375
 - Internet Edge*, 333

- Meraki*, 359–362
- MGMT*, 338
- MPLS Edge*, 333
- overall design architecture*, 328–330
- overview*, 327–328
- QoS classifications and markings*, 352–353
- receiving routes from remote sites*, 348–351
- retail sites*, 359
- SD-WAN*, 343–347
- SD-WAN, VPN*, 350
- SD-WAN headends*, 334–335
- secure Internet access*, 353–355
- SGT*, 359–369
- SIG*, 353–355
- TLOC tunnel groups/color*, 347
- trusted firewalls*, 335–337
- TrustSec*, 359–362
- untrusted firewalls*, 333–334
- transportation use case
 - campus networks*, 425–426
 - campus networks*,
 - macrosegmentation*, 429–432
 - campus networks*,
 - microsegmentation*, 432–433
 - campus networks, transportation-specific entities*, 426–429
 - data centers*, 422–425
 - data centers*,
 - macrosegmentation*, 431–432
 - Dynamic Tunnel Control Policy*, 437–440
 - macrosegmentation*, 429–432
 - microsegmentation*, 432–433
 - overall design architecture*, 419–421
 - overview*, 419
 - SDA, bus SGT*, 432–433
 - SDA, data center services*, 433–437
 - SDA, macrosegmentation*, 429–431
 - SD-WAN centralized policies*, 437–440
 - virtual networks, mapping*, 431–432
 - virtual networks, SDA*, 431
 - WAN, macrosegmentation*, 431–432
 - VN-to-VRF automated mapping federation, SDA-ACI*, 93–95
 - VN-to-VRF manual mapping federation, SDA-ACI*, 96
 - user VPC, route programming*, 142–143
 - users, SDA security*, 256–257

V

- VGW (Virtual Private Gateways)**,
 - IPsec tunnels with VGW*, 138–139
- virtual networks**
 - financial use case*, 312
 - handoffs with BN*, 23–25
 - transportation use case*
 - mapping virtual networks*, 431–432
 - SDA*, 431
- viruses, cloud computing**, 249–250
- visibility, cloud computing**, 250
- VLAN (Virtual LAN)**, *retail use case*
 - backstage sites*, 357–358
 - retail sites*, 359
- VMM (Virtual Machine Monitors)**
 - ACI integration*, 405–413
 - Kubernetes architectures*, 410–413

VN (Virtual Networks)

peering, 144–145

VN-to-VRF

automated mapping federation,
SDA-ACI, 93–95

manual mapping federation,
SDA-ACI, 96

VNG (Virtual Network Gateways),
IPsec tunnels, 143–144

VNID (Vxlan Network Identifiers),
SDA/SD-WAN segmentation, 49–51

VPC (Virtual PortChannels), route
programming, 142–143

VPN (Virtual Private Networks)

DMVPN, 34, 45, 61–62

ACI, 64–65

ACI, *best practices*, 65–66

ACI, *segmentation*, 65

ACI, *SXP Reflectors*, 65–66

automation, 281–282

CNF SD-WAN, 202–204

SDA, *best practices*, 64

SDA, *configuring*, 62–63

SDA, *policy enforcement*, 64

security, 267–270

segmentation, 18–19

Dynamic Tunnel Control Policy,
439–440

MPLS-SR VPN, 386–387

route-leaking, enterprise MPLS,
183–184

SD-WAN VPN, retail use case, 350

VRF (Virtual Routing and
Forwarding), 17–18, 22, 23

ACI, tenant policies, 78

enterprise MPLS

back-to-back VRF, ACI, 178

inter-VRF routing/merging
services, 182–184

offloading selective VRF,
181–182

enterprise MPLS, ACI, 178

FTD, SDA, 23

Inter-Tenant/Inter-VRF

ACI multi-cloud configurations,
148–150

Intra-Tenant/Intra-VRF,

ACI multi-cloud integrations,
146–147

MPLS-SR, ACI integration, 390–392,
393–395

VN-to-VRF automated mapping
federation, SDA-ACI, 93–95

VN-to-VRF manual mapping
federation, SDA-ACI, 96

VRF-Lite, 23, 194–195

VXLAN (Virtual Extensible LAN)

ACI routing/forwarding

ACI-iVXLAN encapsulation, 82

VXLAN overlays, 81–82

SDA and ACI integrations, 31–34

W

WAN (Wide-Area Networks)

CNF/data center connections, 220

MPLS as WAN transport, 171–172

redundancy, 9–10

SD-WAN, 45

ACI, guidelines/limitations, 111

ACI, pairwise integration,
101–110

ACI segmentation, 56

automation, 278–281

Campus Fabric, 25–26

Campus Fabric, default end-
to-end macrosegmentation,
26–27

*Campus Fabric, integrated/
non-integrated SD-WAN
solutions, 27–30*

Catalyst Center integrations, 48

Catalyst SD-WAN Manager, 48

*Catalyst SD-WAN Manager,
APIC integration, 55–56*

*centralized policies, financial
use case, 322–326*

cloud computing, 57–58

Cloud OnRamp, 238–243

*cloud to Internet/External
Destinations Using
On-Prem L3Out
Connectivity, 157–161*

CNF, 200–202

CNF, DMVPN, 202–204

CNF, E-MPLS, 209–212

CNF, P-MPLS, 206–209

CoR, 59–61

*DC and Branch Using
SD-WAN, 104–108*

enterprise MPLS, 170–178

financial use case, 321–326

*Headend Control Policy,
324–326*

*manufacturing use case,
297–298*

MPLS, 56–57

*multi-site DC (ACI-to-ACI
Intersite) connectivity Using
SD-WAN, 106–110*

one-box solutions, 27, 29–30

retail use case, 343–347

*retail use case, SD-WAN
headends, 334–335*

retail use case, VPN, 350

route architectures, 28–30

SDA, 46

SDA, best practices, 51–52

SDA, inline tagging, 50–51

*SDA, one-box SDA (integrated
solutions), 46*

SDA, segmentation, 49–51

*SDA, two-box SDA
(nonintegrated solutions),
48–49*

SDA integrations, 30–31

security, cloud computing, 267

security, DMZ zones, 262–267

security, overlays, 262–267

security, underlays, 262

*security, zone-pair policies,
264–267*

segmentation, 18–19

SIG, 58–59

*Site ID-Based Country Control
Policy, 322–324*

templates, 279–281

*transportation use case,
SD-WAN centralized
policies, 437–440*

two-box solutions, 27, 28

SD-WAN Edge, 57–58

traffic isolation, 17–18

*transportation use case,
macrosegmentation, 431–432*

*wireless architectures, manufacturing
use case, 295–296*

X - Y

**YAML (Yet Another Markup
Language), SDA, 277–278**

Z

zero trust security, 272

**zone-pair policies, SD-WAN security,
264–267**