



CCDE Study Guide

Marwan Al-shawi

CCDE No. 20130066

Cisco Press

FREE SAMPLE CHAPTER



SHARE WITH OTHERS

CCDE Study Guide

Marwan Al-shawi, CCDE No. 20130066

Cisco Press

800 East 96th Street

Indianapolis, Indiana 46240 USA

CCDE Study Guide

Marwan Al-shawi

Copyright© 2016 Pearson Education, Inc.

Cisco Press logo is a trademark of Cisco Systems, Inc.

Published by:

Cisco Press

800 East 96th Street

Indianapolis, IN 46240 USA

All rights reserved. No part of this book may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or by any information storage and retrieval system, without written permission from the publisher, except for the inclusion of brief quotations in a review.

Printed in the United States of America 1 2 3 4 5 6 7 8 9 0

First Printing October 2015

Library of Congress Control Number: 2015949761

ISBN-13: 978-1-58714-461-5

ISBN-10: 1-58714-461-1

Warning and Disclaimer

This book covers various possible design options available while working across multiple places in the network. As part of the evaluation process, the book stays focused on various technical requirements, business requirements, constraints, and associated implications rather than on providing best practice recommendations.

Every effort has been made to make this book as comprehensive and as accurate as possible. The book does not attempt to teach foundational knowledge. Please supplement your learning and fill in gaps in knowledge by reviewing separate technology-specific publications as part of your journey to become a Design Expert.

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

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

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.

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.

Publisher: Paul Boger

Associate Publisher: Dave Dusthimer

Business Operation Manager, Cisco Press: Jan Cornelissen

Executive Editor: Brett Bartow

Managing Editor: Sandra Schroeder

Senior Development Editor: Christopher Cleveland

Project Editor: Mandie Frank

Copy Editor: Keith Cline

Technical Editors: Andre Laurent, Denise Fishburne

Editorial Assistant: Vanessa Evans

Designer: Mark Shirar

Composition: codeMantra

Proofreader: Laura Hernandez




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.

 CCDE, CCENT, Cisco Eos, Cisco HealthPresence, the Cisco logo, Cisco Lumin, Cisco Nexus, Cisco StadiumVision, Cisco TelePresence, Cisco WebEx, DCE, and Welcome to the Human Network are trademarks. Changing the Way We Work, Live, Play, and Learn and Cisco Store are service marks; and Access Registrar, Aironet, AsyncOS, Bringing the Meeting To You, Catalyst, CCDA, CCDP, CCIE, CCFP, CCNA, CCNP, CCSP, CCVP, Cisco, Cisco Certified Internetwork Expert logo, Cisco IOS, Cisco Press, Cisco Systems, Cisco Systems Capital, the Cisco Systems logo, Cisco Unity, Collaboration Without Limitation, EtherFast, EtherSwitch, Event Center, Fast Step, Follow Me Browsing, FormShare, GigaDrive, HomeLink, Internet Quotient, IOS, iPhone, iQuick Study, IronPort, the IronPort logo, LightStream, Linksys, MediaTone, MeetingPlace, MeetingPlace Chime Sound, MGX, Networkers, Networking Academy, Network Registrar, PCNow, PIX, PowerPanels, ProConnect, ScriptShare, SenderBase, SMARTnet, Spectrum Expert, StackWise, The Fastest Way to Increase Your Internet Quotient, TransPath, WebEx, and the WebEx logo are registered trademarks of Cisco Systems, Inc. and/or its affiliates in the United States and certain other countries.

All other trademarks mentioned in this document or website 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. (0812R)

About the Author

Marwan Al-shawi, CCDE No. 20130066, is a lead design with British Telecom Global Services. He helps large-scale enterprise customers to select the right technology solutions for their business needs and provides technical consultancy for various types of network designs and architectures. Marwan has been in the networking industry for more than 12 years and has been involved in architecting, designing, and implementing various large-scale networks, some of which are global service provider-grade networks. Marwan has also worked as a technical consultant with Dimension Data Australia, a Cisco Global Alliance Partner; network architect with IBM Australia global technology services; and other Cisco partners and IT solution providers. He holds a Master of Science degree in internetworking from the University of Technology, Sydney. Marwan also holds other certifications such as Cloud Architect Expert (EMCCAE), Cisco Certified Design Professional (CCDP), Cisco Certified Network Professional – Voice (CCNP Voice), and Microsoft Certified Systems Engineer (MCSE). Marwan was selected as a Cisco Designated VIP by the Cisco Support Community (CSC) (official Cisco Systems forums) in 2012, and by the Solutions and Architectures subcommunity in 2014. In addition, in 2015, Marwan was selected as a member of the Cisco Champions program.

About the Technical Reviewers

Andre Laurent, CCDE No.20120024, CCIE NO.21840 (RS, SP, Security) is a Technical Solutions Architect representing Cisco's Commercial West Area. He has been in IT engineering and consulting for his entire career. Andre is a triple CCIE and CCDE, joint certifications held by fewer than 50 people in the world. Outside of his own personal development, Andre has an equal passion about developing others and assisting them with the certification process. Andre is recognized by the Cisco Learning Network as a subject matter expert in the areas of routing, switching, security, and design. Although he wears a Cisco badge, Andre takes a neutral approach in helping clients establish a long-term business and technology vision covering necessary strategy, execution, and metrics for measuring impact. He spends a great deal of time conducting customer workshops, developing design blueprints, and creating reference models to assist customers in achieving quantified and realized business benefits. Andre has built reference architectures in numerous industries such as banking, retail, utilities and aerospace. He also works closely with some of the largest gaming and hospitality companies in the world.

Denise "Fish" Fishburne, CCDE No.20090014, CCIE No.2639, is an engineer and team lead with the Customer Proof of Concept Lab (CPOC) in North Carolina. Fish is a geek who adores learning and passing it on. She works on many technologies in the CPOC, but her primary technical strength seems, however, to be troubleshooting. Fish has been with Cisco since 1996, CPOC since 2001, and has been a regular speaker at Networkers/Cisco Live since 2006.

Dedication

I would like to dedicate this book to my wonderful mother for her continued support, love, encouragement, guidance, and wisdom.

And most importantly, I would like to thank God for all the blessings in my life.

—*Marwan*

Acknowledgments

A special and big thank you goes to the Pearson-Cisco Press team, especially Brett Bartow and Chris Cleveland, for the support and suggestions that made this book possible. It is a pleasure to work with you. I couldn't have asked for a finer team

I'd like to give special recognition to Andre Laurent for providing his expert technical knowledge and experience as a technical editor of this book. Andre's suggestions and feedback from his real-world practical experience as a technical solution architect with Cisco helped to shape and optimize the quality of the content in multiple areas.

I also want to give special recognition to Denise Fishburne for her valuable contribution and input. As usual, she's not afraid to tell you when you're wrong. In addition, the technical accuracy and insight regarding the technologies and design considerations provided by Denise from her long and extensive experience with Cisco POC helped to enhance the accuracy and quality of the content across multiple sections.

In addition, special a special thank you to Elaine Lopes (CCDE and CCAr Program Manager) for her continued encouragement ever since this book was only an idea.

Also, a special and big thank you to the following experts for their valuable time and their expert perspectives about some chapters and sections in this book, which added a significant value to optimize the contents:

Russ White, Orhan Ergun, Diptanshu Singh, and Ivan Pepelnjak.

Contents at a Glance

Introduction xx

Part I Business-Driven Strategic Network Design 1

Chapter 1 Network Design Requirements: Analysis and Design Principles 3

Part II Next Generation - Converged Enterprise Network Architectures 31

Chapter 2 Enterprise Layer 2 and Layer 3 Design 35

Chapter 3 Enterprise Campus Architecture Design 119

Chapter 4 Enterprise Edge Architecture Design 143

Part III Service Provider Networks Design and Architectures 203

Chapter 5 Service Provider Network Architecture Design 205

Chapter 6 Service Provider MPLS VPN Services Design 245

Chapter 7 Multi-AS Service Provider Network Design 329

Part IV Data Center Networks Design 361

Chapter 8 Data Center Network Design 363

Part V High Availability 429

Chapter 9 Network High-Availability Design 431

Part VI Other Network Technologies and Services 473

Chapter 10 Design of Other Network Technologies and Services 475

Appendix References 577

Contents

	Introduction	xx
Part I	Business-Driven Strategic Network Design	1
Chapter 1	Network Design Requirements: Analysis and Design Principles	3
	Design Scope	4
	Business Requirements	5
	Business Continuity	6
	Elasticity to Support the Strategic Business Trends	7
	IT as a “Business Innovation” Enabler	8
	The Nature of the Business	9
	Business Priorities	9
	Functional Requirements	9
	Technical Requirements	10
	Application Requirements	10
	Design Constraints	12
	Crafting the Design Requirements	13
	Planning	16
	Decision Tree	17
	Decision Matrix	17
	Planning Approaches	18
	Strategic Balance	18
	Network Design Principles	19
	Reliability and Resiliency	19
	Modularity	20
	Reliable and Manageable Scalability	21
	Fault Isolation and Simplicity	22
	Hierarchy	23
	Responsiveness	25
	Holistic Design Approach	25
	Physical Layout Considerations	26
	No Gold Plating	29
	Summary	29
Part II	Next Generation - Converged Enterprise Network Architectures	31
Chapter 2	Enterprise Layer 2 and Layer 3 Design	35
	Enterprise Layer 2 LAN Design Considerations	35
	Spanning Tree Protocol	36
	VLANs and Trunking	37

Link Aggregation	37
First Hop Redundancy Protocol and Spanning Tree	38
Enterprise Layer 2 LAN Common Design Options	40
<i>Layer 2 Design Models: STP Based (Classical Model)</i>	40
<i>Layer 2 Design Model: Switch Clustering Based (Virtual Switch)</i>	41
<i>Layer 2 Design Model: Daisy-Chained Access Switches</i>	42
Layer 2 LAN Design Recommendations	43
Enterprise Layer 3 Routing Design Considerations	43
IP Routing and Forwarding Concept Review	43
Link-State Routing Protocol Design Considerations	45
<i>Link-State over Hub-and-Spoke Topology</i>	45
<i>Link-State over Full-Mesh Topology</i>	48
OSPF Area Types	49
OSPF Versus IS-IS	53
<i>Further Reading</i>	53
EIGRP Design Considerations	54
<i>EIGRP: Hub and Spoke</i>	55
<i>EIGRP Stub Route Leaking: Hub-and-Spoke Topology</i>	56
<i>EIGRP: Ring Topology</i>	58
<i>EIGRP: Full-Mesh Topology</i>	58
<i>EIGRP Route Propagation Considerations</i>	59
<i>Further Reading</i>	60
Hiding Topology and Reachability Information Design Considerations	60
IGP Flooding Domains Design Considerations	62
<i>Link-State Flooding Domain Structure</i>	63
<i>EIGRP Flooding Domains Structure</i>	69
<i>Routing Domain Logical Separation</i>	70
Route Summarization	76
<i>Summary Black Holes</i>	78
<i>Suboptimal Routing</i>	80
IGP Traffic Engineering and Path Selection: Summary	83
OSPF	83
IS-IS	84
EIGRP	84
Summary of IGP Characteristics	84
BGP Design Considerations	85
Interdomain Routing	86
BGP Attributes and Path Selection	88

BGP as the Enterprise Core Routing Protocol	89
<i>Enterprise Core Routing Design Models with BGP</i>	90
<i>BGP Shortest Path over the Enterprise Core</i>	94
BGP Scalability Design Options and Considerations	96
<i>BGP Route Reflection</i>	96
<i>Update Grouping</i>	102
<i>BGP Confederation</i>	103
<i>Confederation Versus Route Reflection</i>	105
Further Reading	106
Route Redistribution Design Considerations	107
Single Redistribution Boundary Point	107
Multiple Redistribution Boundary Points	108
<i>Metric Transformation</i>	109
<i>Administrative Distance</i>	110
Route Filtering Versus Route Tagging with Filtering	110
Enterprise Routing Design Recommendations	114
Determining Which Routing Protocol to Use	115
Summary	117
Chapter 3 Enterprise Campus Architecture Design	119
Enterprise Campus: Hierarchical Design Models	119
Three-Tier Model	120
Two-Tier Model	120
Enterprise Campus: Modularity	121
When Is the Core Block Required?	122
Access-Distribution Design Model	123
Enterprise Campus: Layer 3 Routing Design Considerations	126
EIGRP Versus Link State as a Campus IGP	128
Enterprise Campus Network Virtualization	129
Drivers to Consider Network Virtualization	129
Network Virtualization Design Elements	131
Enterprise Network Virtualization Deployment Models	132
<i>Device Virtualization</i>	133
<i>Path Isolation</i>	133
<i>Service Virtualization</i>	136
Summary	141
Further Reading	141

Chapter 4	Enterprise Edge Architecture Design	143
	Enterprise WAN Module	143
	WAN Transports: Overview	144
	Modern WAN Transports (Layer 2 Versus Layer 3)	145
	<i>Layer 2 MPLS-Based WAN</i>	146
	<i>Layer 3 MPLS-Based WAN</i>	148
	Internet as WAN Transport	151
	<i>Internet as WAN Transport Advantages and Limitations</i>	152
	WAN Transport Models Comparison	153
	WAN Module Design Options and Considerations	155
	<i>Design Hierarchy of the Enterprise WAN Module</i>	155
	<i>WAN Module Access to Aggregation Layer Design Options</i>	156
	<i>WAN Edge Connectivity Design Options</i>	158
	<i>Single WAN Provider Versus Dual Providers</i>	160
	Remote Site (Branch) WAN Design Considerations	161
	<i>Internet as WAN Transport (DMVPN Based)</i>	164
	Enterprise WAN Module Design Options	166
	<i>Option 1: Small to Medium</i>	166
	<i>Option 2: Medium to Large</i>	167
	<i>Option 3: Large to Very Large</i>	169
	WAN Virtualization and Overlays Design Considerations and Techniques	170
	WAN Virtualization	172
	<i>Over-the-Top WAN Virtualization Design Options (Service Provider Coordinated/Dependent)</i>	174
	<i>Over-the-Top WAN Virtualization Design Options (Service Provider Independent)</i>	176
	<i>Comparison of Enterprise WAN Transport Virtualization Techniques</i>	181
	<i>WAN Virtualization Design Options Decision Tree</i>	183
	Enterprise WAN Migration to MPLS VPN Considerations	184
	Migrating from Legacy WAN to MPLS L3VPN WAN Scenario	184
	Enterprise Internet Edge Design Considerations	188
	Internet Edge Architecture Overview	188
	Enterprise Multihomed Internet Design Considerations	190
	<i>Multihoming Design Concept and Drivers</i>	190
	<i>BGP over Multihomed Internet Edge Planning Recommendations</i>	192

	<i>BGP Policy Control Attributes for Multihoming</i>	192
	<i>Common Internet Multihoming Traffic Engineering Techniques over BGP</i>	194
	<i>Scenario 1: Active-Standby</i>	194
	Asymmetrical Routing with Multihoming (Issue and Solution)	199
	Summary	202
Part III	Service Provider Networks Design and Architectures	203
Chapter 5	Service Provider Network Architecture Design	205
	Service Provider Network Architecture Building Blocks	207
	Point of Presence	208
	Service Provider Network Core	211
	Service Provider Control Plane Logical Architectures	212
	IGP in Service Provider Networks	212
	BGP in Service Provider Networks	213
	<i>BGP Route Aggregation (ISP Perspective)</i>	213
	<i>Hot- and Cold-Potato Routing (SP Perspective)</i>	217
	Multiprotocol Label Switching	223
	<i>MPLS Label-Switched Path</i>	225
	<i>MPLS Deployment Modes</i>	225
	Multiprotocol BGP	226
	MPLS Traffic Engineering	227
	Business and Technical Drivers	227
	MPLS-TE Planning	231
	<i>MPLS-TE Strategic Planning Approach</i>	231
	<i>MPLS-TE Tactical Planning Approach</i>	232
	MPLS-TE Design Considerations	233
	<i>Constrained Path Calculation</i>	234
	<i>MPLS-TE Tunnel Placement</i>	237
	<i>Routing Domains</i>	239
	<i>Forwarding Traffic Via the TE Tunnel</i>	241
	Summary	243
	Further Reading	244
Chapter 6	Service Provider MPLS VPN Services Design	245
	MPLS VPN (L3VPN)	245
	MPLS L3VPN Architecture Components	246
	<i>L3VPN Control Plane Components</i>	248

<i>L3VPN Forwarding Plane</i>	251
L3VPN Design Considerations	253
<i>Load Sharing for Multihomed L3VPN CE</i>	253
<i>MPLS L3VPN Topologies</i>	254
<i>MP-BGP VPN Internet Routing</i>	262
PE-CE L3VPN Routing Design	264
<i>PE-CE Routing Design Considerations</i>	265
<i>PE-CE Routing Protocol Selection</i>	266
<i>PE-CE Design Options and Recommendations</i>	266
Layer 2 MPLS VPN (L2VPN)	282
IP NGN Carrier Ethernet	284
Virtual Private Wire Service Design Considerations	287
<i>Transport Models</i>	287
<i>VPWS Control Plane</i>	289
Virtual Private LAN Service Design Considerations	291
<i>VPLS Architecture Building Blocks</i>	292
<i>VPLS Functional Components</i>	292
<i>Virtual Switching Instance</i>	293
<i>VPLS Control Plane</i>	293
<i>VPLS Design Models</i>	294
<i>Ethernet Access Model</i>	298
<i>MPLS Access Model</i>	299
<i>H-VPLS with Provider Backbone Bridging</i>	301
<i>EVPN Design Model (Next-Generation MPLS L2VPN)</i>	307
<i>EVPN BGP Routes and Extended Communities</i>	311
Final Thoughts: L2VPN Business Value and Direction	314
Service Provider Control Plane Scalability	315
IGP Scalability Considerations	316
Route Reflection Design Options in SP Networks	318
<i>Provider Routers as RRs for MPLS-VPN</i>	319
<i>Separate RR for MPLS-VPN and IPv4/v6</i>	319
<i>Separate RR per Service (MPLS-VPN and IPv4/v6)</i>	320
<i>Hierarchical RR</i>	321
<i>Partitioned MPLS-VPN RR</i>	323
Hierarchical LSP (Unified MPLS)	325
Summary	327
Further Reading	327

Chapter 7	Multi-AS Service Provider Network Design	329
	Inter-AS Design Options and Considerations	330
	Inter-AS Option A: Back-to-Back VRF (VRF-to-VRF)	330
	Inter-AS Option B: ASBR to ASBR with MP-eBGP Approach	331
	<i>Option B-1: Next-Hop-Self Approach</i>	331
	<i>Option B-2: Redistribute Connected Approach</i>	332
	<i>Option B-3: Multihop MP-eBGP Approach</i>	334
	Inter-AS Option C: Multihop MP-eBGP Between RR	335
	Inter-AS Option D	335
	Inter-AS IPv6 VPN	336
	Inter-AS MPLS-TE	337
	Inter-AS L2VPN	338
	Inter-AS QoS	343
	Comparison of Inter-AS Connectivity Options	344
	Carrier Supporting Carrier	346
	Non-MPLS Customer over MPLS VPN Carrier	346
	MPLS Customer over MPLS VPN Carrier	347
	MPLS VPN Customer over MPLS VPN Carrier	348
	MPLS VPN Customer over MPLS Carrier	348
	MPLS VPN Customer over IP-Only Carrier	349
	Acquisition of an MPLS-L3VPN Service Provider Design Scenario	353
	Background Information	353
	Design Requirements	353
	Available Interconnection Options	354
	Inter-AS Connectivity Model Selection	355
	Proposed Solution	356
	Network Merger implementation Plan	358
	Summary	358
Part IV	Data Center Networks Design	361
Chapter 8	Data Center Networks Design	363
	Traditional Data Center Network Architecture	364
	STP-Based Data Center Network Architecture	365
	mLAG-Based Data Center Network Architecture	367
	Next-Generation Data Center Network Design	367
	Data Center Virtualization and Cloud-Based Services	
	Overview	368
	Drivers Toward New Fabric-Based Data Center Network	
	Architectures	369

Modern Data Center Network Architectures and Overlays	372
<i>Clos Architecture</i>	374
<i>Clos Transport Protocols</i>	376
<i>MAC-in-MAC</i>	377
<i>MAC-in-IP</i>	380
<i>MPLS Based</i>	383
Comparison of Data Center Network Architectures	387
Data Center Interconnect	389
DCI Building Blocks	392
DCI Connectivity Options	393
<i>Routed DCI</i>	394
<i>Layer 2 DCI</i>	398
<i>Dark Fiber-Based DCI</i>	401
<i>Layer 2 DCI over ME Transport</i>	403
<i>TRILL-FabricPath-Based DCI</i>	404
<i>Overlay Transport Virtualization</i>	406
<i>VxLAN-Based DCI</i>	408
DCI Design Considerations	411
<i>SAN Extension</i>	414
<i>DCI Path Optimization Techniques</i>	417
<i>DNS Based</i>	421
<i>Route Health Injection</i>	422
<i>Locator/ID Separation Protocol</i>	423
Summary	428
Further Reading	428
Part V	High Availability 429
Chapter 9	Network High-Availability Design 431
Fault Tolerance	434
Fate Sharing and Fault Domains	436
Network Resiliency Design Considerations	438
Device-Level Resiliency	441
Protocol-Level Resiliency	443
<i>Network Restoration</i>	444
<i>Network Protection Approach</i>	454
<i>BGP FRR</i>	466
Summary	469
Further Reading	470

Part VI	Other Network Technologies and Services	473
Chapter 10	Design of Other Network Technologies and Services	475
	IPv6 Design Considerations	475
	IPv6 Business and Technical Drivers	476
	IPv6 Addressing Types (Review)	477
	Migration and Integration of IPv4 and IPv6	478
	<i>Discovery Phase</i>	479
	<i>Solution Assessment and Planning</i>	479
	<i>Detailed Design</i>	484
	<i>Deployment, Monitoring, and Optimization</i>	488
	Transition to IPv6: Scenario	488
	<i>Network Requirements Analysis</i>	490
	<i>Design Approach</i>	490
	Further Reading	492
	IP Multicast Design Considerations	492
	Enterprise Multicast Design Options and Considerations	494
	<i>Application Characteristic</i>	494
	<i>Multicast IP Address Mapping into Ethernet MAC Address</i>	494
	<i>Multicast Layer 3 Routing</i>	497
	<i>Multicast BGP</i>	506
	<i>Multicast Source Discovery Protocol</i>	507
	<i>Embedded RP</i>	509
	SP Multicast Design Options and Considerations	510
	<i>MVPN (Draft-Rosen Model)</i>	510
	<i>MVPN - Label Switch Multicast</i>	511
	<i>Next-Generation MVPN</i>	512
	Multicast Resiliency Design Considerations	514
	<i>Anycast RP</i>	514
	<i>Anycast-RP Using PIM</i>	515
	<i>Phantom RP</i>	516
	<i>Live-Live Streaming</i>	517
	<i>First Hop Redundancy Protocol-Aware PIM</i>	519
	Final Thoughts on IP Multicast Design	520
	Further Reading	520
	QoS Design Considerations	521
	QoS High Level Design: Business-Driven Approach	521
	QoS Architecture	523

QoS DiffServ Architecture and Toolset	523
<i>Traffic Classification and Marking</i>	525
<i>Traffic Profiling and Congestion Management</i>	528
<i>Congestion Avoidance (Active Queue Management)</i>	531
<i>Admission Control</i>	531
QoS Design Strategy	532
Enterprise QoS Design Considerations	537
<i>Enterprise Campus</i>	537
<i>Enterprise Edge</i>	538
Service Provider QoS Design	543
<i>Traffic Marking Strategy</i>	543
<i>DiffServ MPLS-TE (DS-TE)</i>	547
Further Reading	549
Network Security Design	550
Network Security Design Fundamentals	551
<i>Top-Down Design</i>	551
<i>Security Policy Considerations</i>	551
<i>Holistic Approach Considerations</i>	552
<i>Divide-and-Conquer Approach</i>	553
<i>Security Triad Principle (Confidentiality, Integrity, and Availability)</i>	555
Network Infrastructure Security Considerations	556
<i>Network Device Level Security</i>	557
<i>Layer 2 Security Considerations</i>	561
<i>Layer 3 Control Plane Security Considerations</i>	563
<i>Remote-Access and Network Overlays (VPN) Security Considerations</i>	564
<i>Network-Based Firewall Considerations</i>	566
Further Reading	568
Network Management	569
Fault, Configuration, Accounting, Performance, and Security	570
Network Management High-Level Design Considerations	571
Multitier Network Management Design	574
Further Reading	576
Summary	576

Icons Used in This Book



Layer 2 Switch



Layer 3 Switch



Modular Layer 3
Switch




Frame-Relay/ATM
WAN Switch



Router



MPLS Router



Layer 2 WAN/SP
Aggregation
Switch



SAN Switch



Router with
IP Tunnel



Firewall



Satellite



Host with
Virtual
Machines



Load Balancer



Fabric Switch



IP Phone



Workstation



Server



Remote or
Regional Site



Radio Tower



Optical Ring



Virtual Machine



Ethernet Link



Legacy Link-Serial,
Frame-Relay, ATM, TDM



Cloud-Routed or
Switched Domain

Introduction

The CCDE certification is a unique certification in the IT and networking industry and is considered one of the most if not the only recognized vendor-neutral network Design Expert level certification. When it comes to design, it is like art: It cannot be taught or covered entirely through a single book or a training course, because each design has different drivers, philosophy, and circumstances that collectively create its unique characteristic. Therefore, this book uses a comparative and analytical approach to help the reader answer the question why with regard to design or technology selections, and to think of how to link the technical design decisions to other influencing factors (technical, nontechnical, or combination of both) to achieve a true and successful business-driven design. In addition, multiple mini design scenarios and illustrations are included in the chapters to explain the concepts, design options, and implications.

This book is the first book to target the CCDE practical exam. Also, It is the first book that consists of diverse design aspects, principles, and options using a business-driven approach for enterprise, service provider, and data center networks.

This book covers the different design principles and topics using the following approach:

- Covers (that is, highlights, discusses, and compares) the different technologies, protocols, design principles, and design options in terms of cost, availability, scalability, performance, flexibility, and so on (along with the strength of the various designs and design concerns where applicable)
- Covers the drivers toward adopting the different technologies and protocols (technical and nontechnical) (whether intended for enterprise or service provider networks depends on the topic and technology)
- Covers the implications of the addition or integration of any element to the overall design, such as adding new applications or integrating two different networks

The design topics covered in this CCDE Study Guide aim to prepare you to be able to

- Analyze and identify various design requirements (business, functional, and application) that can influence design decisions.
- Understand the different design principles and their impact on the organization from a business point of view
- Understand and compare the various network design architectures and the associated implications on various design aspects of applying different Layer 2 and Layer 3 control plane protocols
- Identify and analyze design limitations or issues, and how to optimize them, taking into consideration the technical and nontechnical design requirements and constraints.
- Identify and analyze the implication of adding new services or applications and how to accommodate the design or the design approach to meet the expectations

This book references myriad sources, but presents the material in a manner tailored for conscientious network designers and architects. The material also covers updated standards and features that are found in enterprise and service provider networks. In addition, each chapter contains further reading suggestions pointing you to recommended documents that pertain to the topics covered in each chapter.

Therefore, you can use this book as an all-in-one study guide covering the various networking technologies, protocols, and design options in a business-driven approach. You can expand your study scope and depth of knowledge selectively on certain topics as needed.

Whether you are preparing for the CCDE certification or just want to better understand advanced network design topics, you will benefit from the range of topics covered and the practical business-driven approach used to analyze, compare, and explain these topics.

Who Should Read This Book?

In addition to those who are planning or studying for the CCDE certification, this book is for network engineers, network consultants, network architects, and solutions architects who already have a foundational knowledge of the topics being covered and who would like to train themselves to think more like a design engineer rather than like an implementation engineer.

CCDE Practical Exam Overview

The minimally qualified CCDE must have expert-level knowledge, experience, and skills that cover complex networks design (ideally global-scale networks) by successfully demonstrating the ability to translate business requirements and strategies into functional and technical strategies. In other words, CCDEs are recognized for their expert-level knowledge and skills in network infrastructure design. The deep technical networking knowledge that a CCDE brings ensures that they are well qualified to address the most technically challenging network infrastructure design assignments [1]. Therefore, to test the CCDE candidate skills, knowledge, and expertise, the CCDE practical exam is divided into multiple design scenarios, with each having a different type of network and requirements. In addition, each design scenario is structured of different domains and tasks that CCDE candidates have to complete to demonstrate expert-level abilities when dealing with a full network design lifecycle (gather business requirements, analyze the requirements, develop a design, plan the implementation of the design, and then apply and optimize the design).

Job Tasks

The CCDE exam is designed to cover different use cases, each of which may be integrated in one or multiple design scenarios in the exam. The following are the primary use cases at the time of this writing:

- **Merge or divest networks:** This use case covers the implications and challenges (technical and nontechnical) associated with merging or separating different networks (both enterprise and service provider types of networks). This domain, in particular, can be one of the most challenging use cases for network designers because, most of the time, merging two different networks means integrating two different design philosophies, in which multiple conflicting design concepts can appear. Therefore, at a certain stage, network designers have to bring these two different networks together to work as a one cohesive system, taking into consideration the various design constraints that might exist such as goals for the merged network, security policy compliance, timeframe, cost, the merger constraints, the decision of which services to keep and which ones to divest (and how), how to keep services up and running after the divestiture, what the success criteria is for the merged network, and who is the decision maker.
- **Add technology or application:** This use case covers the impact of adding technology or an application to an existing network. Will anything break as a result of the new addition? In this use case, you must consider the application requirements in terms of traffic pattern, convergence time, delay, and so on across the network. By understanding these requirements, the CCDE candidate as a network designer should be able to make design decisions about fine-tuning the network (design and features such as quality of service [QoS]) to deliver this application with the desired level of experience for the end users.
- **Replace technology:** This use case covers a wide range of options to replace an existing technology to meet certain requirements. It might be a WAN technology, routing protocol, security mechanism, underlying network core technology, or so on. Also, the implications of this new technology or protocol on the current design, such as enhanced scalability or potential to conflict with some of the existing application requirements, require network designers to tailor the network design so that these technologies work together rather than in isolation, so as to reach objectives, such as delivering business applications and services.
- **Scale a network:** This use case covers different types of scalability aspects at different levels, such as physical topology, along with Layer 2 and Layer 3 scalability design considerations. In addition, the challenges associated with the design optimization of an existing network design to offer a higher level of scalability are important issues in this domain. For example, there might be some constraints or specific business requirements that might limit the available options for the network

designer when trying to optimize the current design. Considerations with regard to this use case include the following: Is the growth planned or organic? Are there issues caused by the growth? Should one stop and redesign the network to account for growth? What is the most scalable design model?

Exam Job Domains

In each of the CCDE exam use cases described in the preceding section, as part of each CCDE design scenario the candidate is expected to cover some or all of the following job domains:

- **Analyze:** This domain requires identification of the requirements, constraints, and risks from both business and technology perspectives. In this task, the candidate is expected to perform multiple subtasks, such as the following:
 - Identify the missing information required to produce a design.
 - Integrate and analyze information from multiple sources (for example, from e-mails or from diagrams) to provide the correct answer for any given question.
- **Design:** In this domain, the CCDE candidate is usually expected to provide a suggested design by making design choices and decisions based on the given and analyzed information and requirements in the previous task. Furthermore, one of the realistic and unique aspects of the CCDE exam is that there might be more than one right design option. Also, there might be optimal and suboptimal solutions. This aspect of the exam is based on the CCDE candidate's understanding of the requirements, goals, and constraints in making the most relevant and suitable selection given the options available.
- **Implement and deploy:** This domain contains multiple subtasks, such as the following:
 - Determine the consequences of the implementation of the proposed design.
 - Design implementation, migration, or fallback plans.

Note No command-line interface (CLI) configuration is required on the CCDE practical exam. The general goal behind this point is more about how and where you to apply a network technology or a protocol and the implications associated with it, and how to generate a plan to apply the proposed design.

- **Validate and optimize:** Here the CCDE candidate is required to justify the design choices and decisions in terms of the rationale behind the design's selection. The candidate's justifications should evidence that the selected design is the best available. Justifications are typically driven by technical requirements, business requirements, and functional requirements, considered either in isolation or in combination.

Exam Technologies

As a general rule for the CCDE practical exam technologies, consider the written exam (qualification) blueprint as a reference (see Figure I-1). However, remember that this is a scenario-based design exam, in which you might expect expansion to the technologies covered in the CCDE written blueprint. In other words, consider the blueprint as a foundation and expand upon it to a reasonable extent; it is not necessary to go deeply into technologies that are not used in real-life networks.

L2 Control Plane	<ul style="list-style-type: none"> • STP • Switch Fabric
L3 Control Plane	<ul style="list-style-type: none"> • IGP (IS-IS, OSPF, EIGRP) • BGP (iBGP, eBGP, MP-BGP)
Network Virtualization	<ul style="list-style-type: none"> • L2VPN, L3VPN, VRF • Tunneling (IPSec, GRE, DMVPN)
QoS	<ul style="list-style-type: none"> • DiffServ, Inserv
Network Security	<ul style="list-style-type: none"> • Service Provider and Enterprise network infrastructure security
Network Management	<ul style="list-style-type: none"> • SNMP, Netflow

Figure I-1 Exam Technologies

Note The above technologies include both IP versions (version 4 and 6) as well as unicast and Multicast IP communications.

PPDIOO Approach and the CCDE Job Domains

With regard to IT services, businesses usually aim to reduce total cost of ownership, improve its service availability, enhance user quality of experience, and reduce operational expenses. By adopting a lifecycle approach, organizations can define a set of methodologies and standards to be followed at each stage of the IT network's life. With this approach, there will be a series of phases that all collectively construct a lifecycle. With most lifecycle approaches, the information and findings of each phase can be used to feed and improve the following phase. This ultimately can produce more efficient and cost-effective IT network solutions that offer IT more elasticity to enhance current investments and to adopt new IT technologies and justify their investment cost.

The PPDIOO lifecycle (see Figure I-2) stands for Prepare, Plan, Design, Implement, Operate, and Optimize, which is Cisco's vision of the IT network lifecycle. This vision is

primarily based on the concept that understanding what is supposed to happen at each stage is vital for a company (or architect, designer) to properly use the lifecycle approach and to get the most benefit from it.

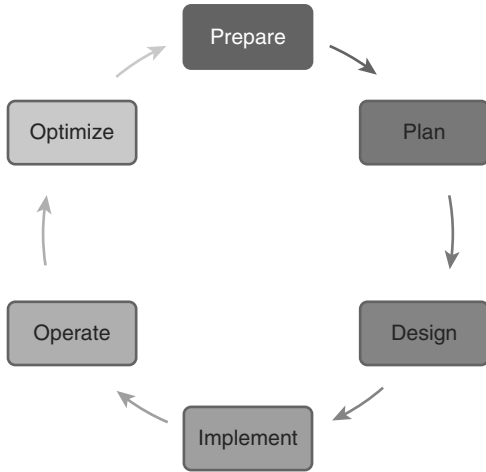


Figure I-2 PPDIOO Lifecycle

Furthermore, this approach offers the flexibility to have a two-way directional relationship between the phases. For instance, during the monitoring phase of the newly designed and implemented network, issues might be discovered that can be fixed by the addition of some features. This is similar to when there are issues related to design limitations. Therefore, each phase can provide reverse feeding, as well, to the previous phase or phases to overcome issues and limitations that appear during the project lifecycle. As a result, this will provide an added flexibility to IT in general and the design process in particular to provide a workable design that can transform the IT network infrastructure to be a business enabler. Figure I-3 illustrates the PIDDO lifecycle with the multidimensional relationship between the lifecycle phases.




Figure I-3 PPDIOO Multidimensional Relationship

PPDIOO and Tasks

In fact, the PPDIOO lifecycle is applicable to the CCDE job domains, just like any other design project:

- The CCDE candidate needs to analyze the provided information (Prepare).
- Use this information to make design choices and decisions (Plan).
- Generate, propose, or suggest a suitable design (Design).
- Apply the selected design (for example, an implementation plan) (Implement).
- Collect feedback or monitor (Operate) for optimization and enhancement (Optimize).

Final Thoughts on the CCDE Practical Exam

Understanding the various domains and tasks expected in each of the exam's design scenarios can help CCDE practical exam candidates shape their study plans. This understanding can also help those who have the opportunity to practice it in their work environment. If they are working on a design and architecture project, they will have a tangible practical feeling and understand how the different stages of the design process can be approached and handled.

How This Book Is Organized

Although this book could be read cover to cover, it is designed to be flexible and allow you to easily move between chapters and sections of chapters to cover just the topics that you need more work with.

This book is organized into six distinct sections.

Part I of the book explains briefly the various design approaches, requirements, and principles, and how they complement each other to achieve a true business-driven design.

- **Chapter 1, “Network Design Requirements: Analysis and Design Principles”**
This chapter covers how the different requirements (business, functional, and application) can influence design decisions and technology selection to achieve a business-driven design. This chapter also examines, when applicable, the foundational design principles that network designers need to consider.

Part II of this book focuses on the enterprise networks, specifically modern (converged) networks. The chapter covers the various design options, considerations, and design implications with regard to the business and other design requirements.

- **Chapter 2, “Enterprise Layer 2 and Layer 3 Design”** This chapter covers different design options and considerations related to Layer 2 and Layer 3 control plane protocols and advanced routing concepts.

- **Chapter 3, “Enterprise Campus Architecture Design”** This chapter covers the design options applicable to modern campus networks. The chapter also covers some of the design options and considerations with regard to network virtualization across the campus network.

- **Chapter 4, “Enterprise Edge Architecture Design”** This chapter covers various design options and considerations with regard to the two primary enterprise edge blocks (WAN edge and Internet edge).

Part III of the book focuses on service provider-grade networks. It covers the various design architectures, technologies, and control protocols, along with the drivers toward adopting the different technologies (technical and nontechnical).

- **Chapter 5, “Service Provider Network Architecture Design”** This chapter covers the various architectural elements that collectively comprise a service provider-grade network at different layers (topological and protocols layers). The chapter also covers the implications of some technical design decisions on the business.
- **Chapter 6, “Service Provider MPLS VPN Services Design”** This chapter covers various options and considerations in MPLS VPN network environments, focusing on L2VPN and L3VPN networks. The chapter also examines different design options and approaches to optimize Layer 3 control plane design scalability for service provider-grade networks.
- **Chapter 7, “Multi-AS Service Provider Network Design”** This chapter focuses on the design options and considerations when interconnecting different networks or routing domains. The chapter examines each design option and then compares them based on various design aspects such as security and scalability.

Part IV of the book focuses on data center networks design for both traditional and modern (virtualized and cloud based) data center networks. This part also covers how to achieve business continuity goals.

- **Chapter 8, “Data Center Networks Design”** This chapter covers various design architectures, concepts, techniques, and protocols that pertain to traditional and modern data center networks. In addition, this chapter analyzes and compares the different design options and considerations, and examines the associated implications of interconnecting dispersed data center networks and how these different technologies and design techniques can facilitate achieving different levels of business continuity.

Part V of this book focuses on the design principles and aspects to achieve the desired levels of operational uptime and resiliency by the business.

- **Chapter 9, “Network High-Availability Design”** This chapter covers the different variables and factors that either solely or collectively influence the overall targeted operational uptime. This chapter also examines the various elements that influence achieving the desired degree of network resiliency and fast convergence.

Part VI of the book focuses on network technologies and services that are not core components of the CCDE practical exam.

- **Chapter 10, “Design of Other Network Technologies and Services”** This chapter briefly explains some design considerations with regard to the following network technologies and services, with a focus on certain design aspects and principles and without going into deep technical detail or explanation: IPv6, multicast, QoS, security, and network management.

Final Words

This book is an excellent self-study resource to learn how to think like a network designer following a business-driven approach. You will learn how to analyze and compare different design options, principles, and protocols based on various design requirements. However, the technical knowledge forms only the foundation to pass the CCDE practical exam. You also want to have a real feeling for gathering business requirements, analyzing the collected information, identifying the gaps, and producing a proposed design or design optimization based on that information. If you believe that any topic in this book is not covered in enough detail, I encourage you to expand your study scope on that topic by using the recommended readings in this book and by using the recommended CCDE study resources available online at Learning@Cisco.com

Enjoy the reading and happy learning.

This page intentionally left blank

Enterprise Campus Architecture Design

A campus network is generally the portion of the enterprise network infrastructure that provides access to network communication services and resources to end users and devices that are spread over a single geographic location. It may be a single building or a group of buildings spread over an extended geographic area. Normally, the enterprise that owns the campus network usually owns the physical wires deployed in the campus. Therefore, network designers typically tend to design the campus portion of the enterprise network to be optimized for the fastest functional architecture that runs on high-speed physical infrastructure (1/10/40/100 Gbps). Moreover, enterprises can also have more than one campus block within the same geographic location, depending on the number of users within the location, business goals, and business nature. When possible, the design of modern converged enterprise campus networks should leverage the following common set of engineering and architectural principles [10]:

- Hierarchy
- Modularity
- Resiliency

Enterprise Campus: Hierarchical Design Models

The hierarchical network design model breaks the complex flat network into multiple smaller and more manageable networks. Each level or tier in the hierarchy is focused on a specific set of roles. This design approach offers network designers a high degree of flexibility to optimize and select the right network hardware, software, and features to perform specific roles for the different network layers.

A typical hierarchical enterprise campus network design includes the following three layers:

- **Core layer:** Provides optimal transport between sites and high-performance routing. Due to the criticality of the core layer, the design principles of the core should provide

an appropriate level of resilience that offers the ability to recover quickly and smoothly after any network failure event with the core block.

- **Distribution layer:** Provides policy-based connectivity and boundary control between the access and core layers.
- **Access layer:** Provides workgroup/user access to the network.

The two primary and common hierarchical design architectures of enterprise campus networks are the three-tier and two-tier layers models.

Three-Tier Model

This design model, illustrated in Figure 3-1, is typically used in large enterprise campus networks, which are constructed of multiple functional distribution layer blocks.

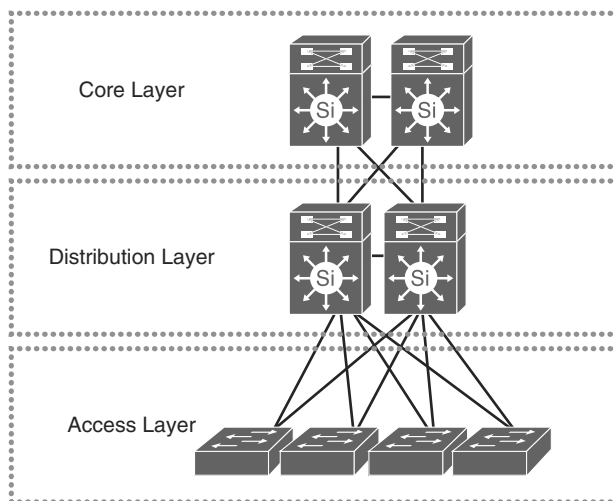


Figure 3-1 *Three-Tier Network Design Model*

Two-Tier Model

This design model, illustrated in Figure 3-2, is more suitable for small to medium-size campus networks (ideally not more than three functional disruption blocks to be interconnected), where the core and distribution functions can be combined into one layer, also known as *collapsed core-distribution architecture*.

Note The term *functional distribution block* refers to any block in the campus network that has its own distribution layer such as user access block, WAN block, or data center block.




Figure 3-2 *Two-Tier Network Design Model*

Enterprise Campus: Modularity

By applying the hierarchical design model across the multiple functional blocks of the enterprise campus network, a more scalable and modular campus architecture (commonly referred to as *building blocks*) can be achieved. This modular enterprise campus architecture offers a high level of design flexibility that makes it more responsive to evolving business needs. As highlighted earlier in this book, modular design makes the network more scalable and manageable by promoting fault domain isolation and more deterministic traffic patterns. As a result, network changes and upgrades can be performed in a controlled and staged manner, allowing greater stability and flexibility in the maintenance and operation of the campus network. Figure 3-3 depicts a typical campus network along with the different functional modules as part of the modular enterprise architecture design.

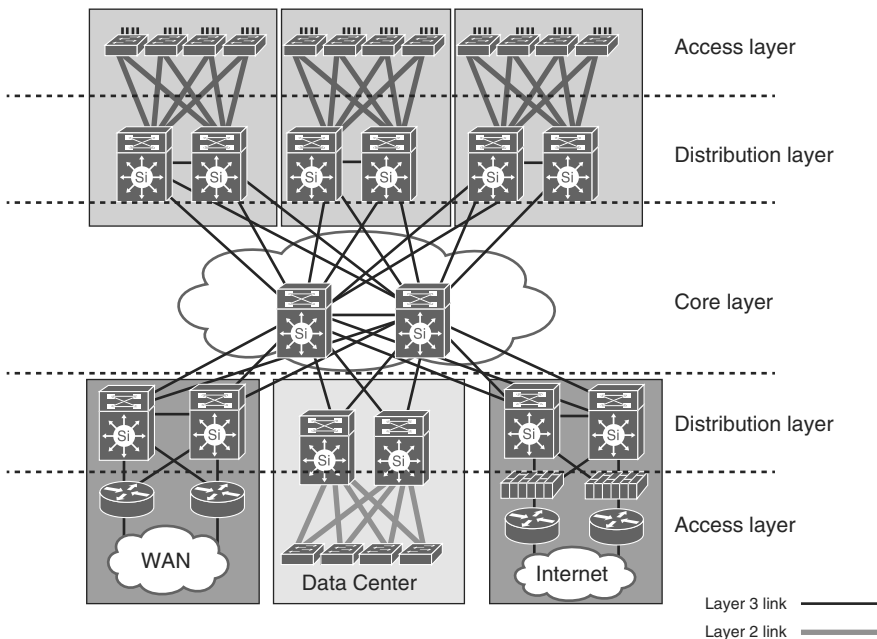


Figure 3-3 *Typical Modular Enterprise Campus Architecture*

Note Within each functional block of the modular enterprise architecture, to achieve the optimal structured design, you should apply the same hierarchal network design principle.

When Is the Core Block Required?

A separate core provides the capability to scale the size of the enterprise campus network in a structured fashion that minimizes overall complexity when the size of the network grows (multiple campus distribution blocks) and the number of interconnections tying the multiple enterprise campus functional blocks increases significantly (typically leads to physical and control plane complexities), as exemplified in Figure 3-4. In other words, not every design requires a separate core.




Figure 3-4 Network Connectivity Without Core Versus With Core

Besides the previously mentioned technical considerations, as a network designer you should always aim to provide a business-driven network design with a future vision based on the principle “build today with tomorrow in mind.” Taking this principle into account, one of the primary influencing factors with regard to selecting two-tier versus three-tier network architecture is the type of site or network (remote branch, regional HQ, secondary or main campus), which will help you, to a certain extent, identify the nature of the site and its potential future scale (from a network design point of view). For instance, it is rare that a typical (small to medium-size) remote site requires a three-tier architecture even when future growth is considered. In contrast, a regional HQ site or a secondary campus network of an enterprise can have a high potential to grow significantly in size (number of users and number of distribution blocks). Therefore, a core layer or three-tier architecture can be a feasible option here. This is from a hypothetical design point of view; the actual answer must always align with the business goals and plans (for example if the enterprise is planning to merge or acquire any new business); it can also derive from the projected percentage of the yearly organic business growth. Again, as a network designer, you can decide based on the current size and the projected growth, taking into account the type of the targeted site, business nature, priorities, and design constraints such as cost. For example, if the business priority is to expand

without spending extra on buying additional network hardware platforms (reduce capital expenditure [capex]), in this case the cost savings is going to be a design constraint and a business priority, and the network designer in this type of scenario must find an alternative design solution such as the collapsed architecture (two-tier model) even though technically it might not be the optimal solution.

That being said, sometimes (when possible) you need to gain the support from the business first, to drive the design in the right direction. By highlighting and explaining to the IT leaders of the organization the extra cost and challenges of operating a network that was either not designed optimally with regard to their projected business expansion plans, or the network was designed for yesterday's requirements and it will not be capable enough to handle today's requirements. Consequently, this may help to influence the business decision as the additional cost needed to consider three-tier architecture will be justified to the business in this case (long-term operating expenditure [opex] versus short-term capex). In other words, sometimes businesses focus only on the solution capex without considering that opex can probably cost them more on the long run if the solution was not architected and designed properly to meet their current and future requirements

Access-Distribution Design Model

Chapter 2, “Enterprise Layer 2 and Layer 3 Design,” discussed different Layer 2 design models that are applicable to the campus LAN design, in particular to the access-distribution layer. Technically, each design model has different design attributes. Therefore, network designers must understand the characteristics of each design model to be able to choose and apply the most feasible model based on the design requirements.

The list that follows describes the three primary and common design models for the access layer to distribution layer connectivity. The main difference between these design models is where the Layer 2 and Layer 3 boundary is placed and how and where Layer 3 gateway services are handled:

- **Classical multitier STP based:** This model is the classical or traditional way of connecting access to the distribution layer in the campus network. In this model, the access layer switches usually operate in Layer 2 mode only, and the distribution layer switches operate in Layer 2 and Layer 3 modes. As discussed earlier in this book, the primary limitation of this design model is the reliance on Spanning Tree Protocol (STP) and First Hop Redundancy Protocol (FHRP). For more information, see Chapter 2.
- **Routed access:** In this design model, access layer switches act as Layer 3 routing nodes, providing both Layer 2 and Layer 3 forwarding. In other words, the demarcation point between Layer 2 and Layer 3 is moved from the distribution layer to the access layer. Based on that, the Layer 2 trunk links from access to distribution are replaced with Layer 3 point-to-point routed links, as illustrated in Figure 3-5.




Figure 3-5 *Routed Access Layer*

The routed access design model has several advantages compared to the multitier classical STP-based access-distribution design model, including the following:

- Simpler and easier to troubleshoot, you can use a standard routing troubleshooting techniques, and you will have fewer protocols to manage and troubleshoot across the network
- Eliminate the reliance on STP and FHRP and rely on the equal-cost multipath (EMCP) of the used routing protocol to utilize all the available uplinks, which can increase the overall network performance
- Minimize convergence time during a link or node failure

Note The routed access design model does not support spanning Layer 2 VLANs across multiple access switches, and this might not be a good choice for some networks. Although expanding Layer 2 over routed infrastructure is achievable using other different overlay technologies, this might add complexity to the design, or the required features may not be supported with the existing platforms for the access or distribution layer switches.

- **Switch clustering:** As discussed in Chapter 2, this design model provides the simplest and most flexible design compared to the other models discussed already. As illustrated in Figure 3-6, by introducing the switch clustering concept across the different functional modules of the enterprise campus architecture, network designers can simplify and enhance the design to a large degree. This offers a higher level of node and path resiliency, along with significantly optimized network convergence time.

The left side of Figure 3-6 represents the physical connectivity, and the right side shows the logical view of this architecture, which is based on the switch clustering design model across the entire modular campus network.

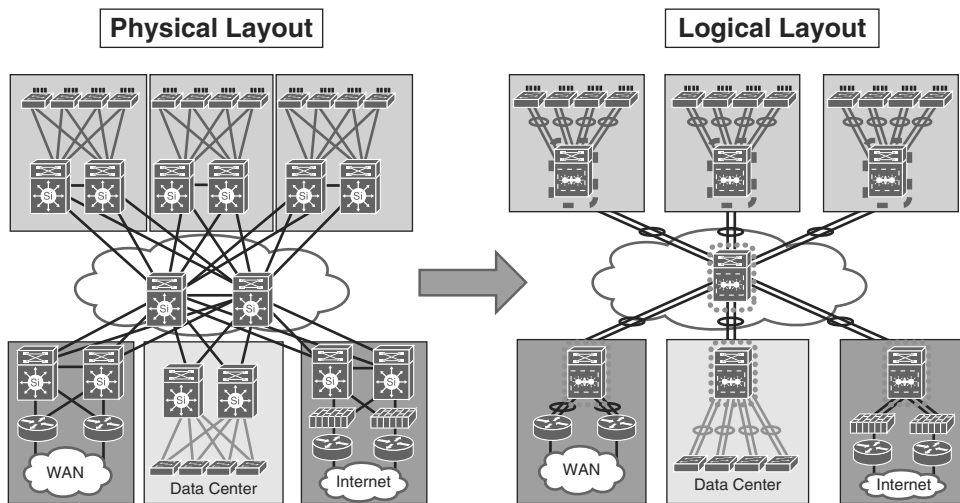


Figure 3-6 Switch Clustering Concept

Table 3-1 compares the different access-distribution connectivity design models from different design angles.

Table 3-1 Comparing Access-Distribution Connectivity Models

	Multitier STP Based	Routed Access	Switch Clustering
Design flexibility	Limited (topology dependent)	Limited (For example, spanning Layer 2 over different access switches requires an overlay technology)	Flexible
Scalability	Supports scale up and limited scale out (topology dependent)	Supports both scale up and scale out	Scale up and limited scale out (typically limited to 2 distribution switches per cluster)
Layer 3 gateway services	Distribution layer (FHRP based)	Access layer (Layer 3 routing based)	Distribution layer (may or may not require FHRP*)
Multichassis link aggregation (mLAG)	Not supported	Not supported (instead relies on Layer 3 ECMP)	Supported

Table 3-1 *continued*

	Multitier STP Based	Routed Access	Switch Clustering
Access-to-distribution convergence time	Dependent on STP and FHRP timers (relatively slow)	Interior Gateway Protocol (IGP) dependent, commonly fast	Fast
Operational complexity	Complex (multiple control protocols to deal with [for example, STP, FHRP])	Moderate (Advanced routing design expertise may be required)	Simple

* Some switch clustering technologies, such as Cisco Nexus vPC, use FHRP (Hot Standby Router Protocol [HSRP]). However, from a forwarding plane point of view, both upstream switches (vPC peers) do forward traffic, unlike the classical behavior, which is based on active-standby.

Note All the design models discussed in this section are valid design options. However, the design choice must be driven by the requirements and design constraints, such as cost, which can influence which option you can select. For example, an access switch with Layer 3 capabilities is more expensive than a switch with Layer 2 capabilities only. This factor will be a valid tiebreaker if cost is a concern from the perspective of business requirements.

Enterprise Campus: Layer 3 Routing Design Considerations

The hierarchal enterprise campus architecture can facilitate achieving more structured hierarchal Layer 3 routing design, which is the key to achieving routing scalability in large networks. This reduces, to a large extent, the number of Layer 3 nodes and adjacencies in any given routing domain within each tier of the hierarchal enterprise campus network [27].

In a typical hierarchal enterprise campus network, the distribution block (layer) is considered the demarcation point between Layer 2 and Layer 3 domains. This is where Layer 3 uplinks participate in the campus core routing, using either an interior routing protocol (IGP) or Border Gateway Protocol (BGP), which can help to interconnect multiple campus distribution blocks together for end-to-end IP connectivity.

By contrast, with the routed access design model, Layer 3 routing is extended to the access layer switches. Consequently, the selection of the routing protocol is important for a redundant and reliable IP/routing reachability within the campus, considering scalability and the ability of the network to grow with minimal changes and impact to the

network and routing design. All the Layer 3 routing design considerations discussed in previous chapters must be considered when applying any routing protocol to a campus LAN. Figure 3-7 illustrates a typical ideal routing design that aligns the IGP design (Open Shortest Path First [OSPF]) with the enterprise campus hierarchical architecture, along with the different functional modules.

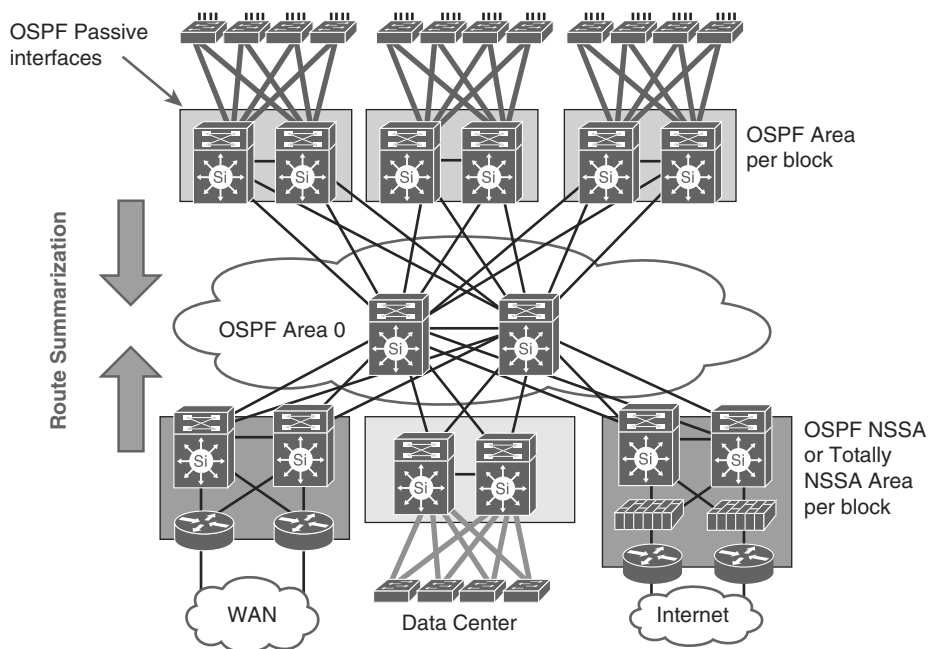


Figure 3-7 Campus Network: Layer 3 Routing

Note In the preceding example, the data center and other modules coexist within the enterprise campus to illustrate a generic IGP design over multiple modules interconnected over one core infrastructure (typical large campus design). However, in some designs, the data center is interconnected over a WAN or dedicated fiber links with the campus network. Also, the other blocks, such as Internet and WAN blocks, might be coresident at the data center's physical location as well. In this case, you can use the same IGP design concept, and you can treat the WAN interconnect as an external link, as illustrated in Figure 3-8. Also, the enterprise core routing (that is, OSPF backbone area) can be extended over the WAN. In other words, all the concepts of IGP and BGP design discussed earlier in this book have to be considered to make the right design decision. (There is no single standard design that you can use in different scenarios.)

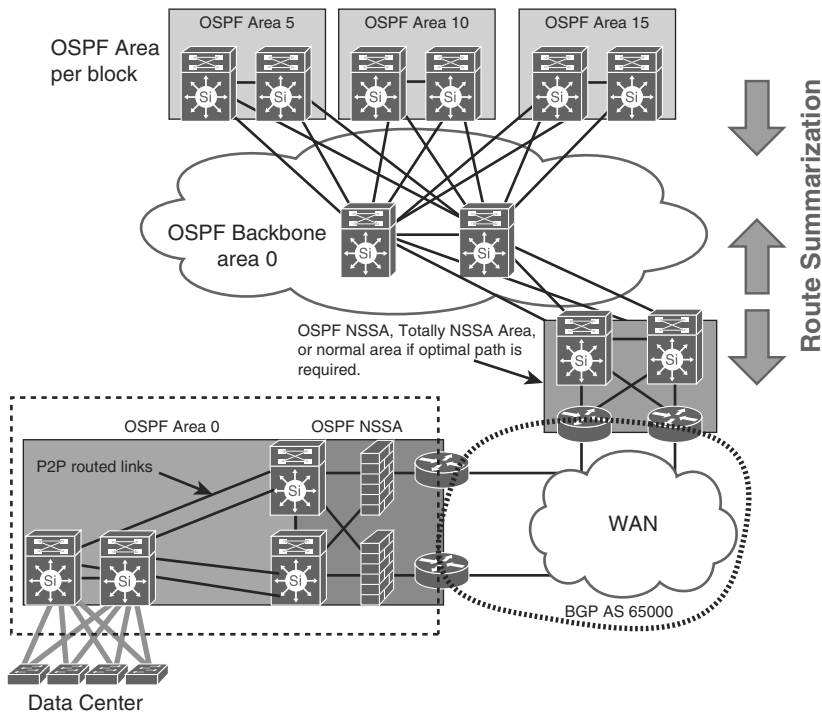


Figure 3-8 *Campus Network: Layer 3 Design with WAN Core*

EIGRP Versus Link State as a Campus IGP

As discussed in Chapter 2, each protocol has its own characteristics, especially when applied to different network topologies. For example, Enhanced Interior Gateway Routing Protocol (EIGRP) offers a more flexible, scalable, and easier-to-control design over “hub-and-spoke” topology compared to link state. In addition, although EIGRP is considered more flexible on multitiered network topologies such as three-tier campus architecture, link-state routing protocols have still proven to be powerful, scalable, and reliable protocols in this type of network, especially OSPF, which is one of the most commonly implemented protocols used in campus networks. Furthermore, in large-scale campus networks, if EIGRP is not designed properly with regard to information hiding and EIGRP query scope containment (discussed in Chapter 2), any topology change may lead to a large floods of EIGRP queries. In addition, the network will be more prone to EIGRP stuck-in-active (SIA) impacts, such as a longer time to converge following a failure event and as a SIA timer puts an upper boundary on convergence times.

Consequently, each design has its own requirements, priorities, and constraints; and network designers must evaluate the design scenario and balance between the technical (protocol characteristics) and nontechnical (business priorities, future plans, staff knowledge, and so on) aspects when making design decisions.

Table 3-2 provides a summarized comparison between the two common and primary IGPs (algorithms) used in large-scale hierarchal enterprise campus networks.

Table 3-2 *Link State Versus EIGRP in the Campus*

Design Consideration	EIGRP (DUAL)	Link State (Dijkstra)
Architecture flexibility	High (natively supports multitier architectures with routes summarization)	High, with limitations (The more tiers the network has, the less flexible the design can be.)
Scalability	High	High
Convergence time (protocol level)*	Fast (ideally with route summarization)	Fast (ideally with topology hiding, route summarization, and timers tuning)
MPLS-TE support	No	Yes

* This design aspect is covered in more detail in Chapter 9, “Network High-Availability Design.”

Enterprise Campus Network Virtualization

Virtualization in IT generally refers to the concept of having two or more instances of a system component or function such as operating system, network services, control plane, or applications. Typically, these instances are represented in a logical virtualized manner instead of being physical.

Virtualization can generally be classified into two primary models:

- **Many to one:** In this model, multiple physical resources appear as a single logical unit. The classical example of many-to-one virtualization is the switch clustering concept discussed earlier. Also, firewall clustering, and FHRP with a single virtual IP (VIP) that front ends a pair of physical upstream network nodes (switches or routers) can be considered as other examples of the many-to-one virtualization model.
- **One to many:** In this model, a single physical resource can appear as many logical units, such as virtualizing an x86 server, where the software (hypervisor) hosts multiple virtual machines (VMs) to run on the same physical server. The concept of network function virtualization (NFV) can also be considered as a one-to-many system virtualization model.

Drivers to Consider Network Virtualization

To meet the current expectations of business and IT leaders, a more responsive IT infrastructure is required. Therefore, network infrastructures need to move from the classical architecture (that is, based on providing basic interconnectivity between different siloed departments within the enterprise network) into a more flexible, resilient, and adaptive architecture that can support and accelerate business initiatives and remove

inefficiencies. The IT and the network infrastructure will become like a service delivery business unit that can quickly adopt and deliver services. In other words, it will become a “business enabler.” This is why network virtualization is considered one of the primary principles that enables IT infrastructures to become more dynamic and responsive to the new and the rapidly changing requirements of today’s enterprises.

The following are the primary drivers of modern enterprise networks, which can motivate enterprise businesses to adopt the concept of network virtualization:

- **Cost efficiency and design flexibility:** Network virtualization provides a level of abstraction from the physical network infrastructure that can offer cost-effective network designs along with a higher degree of design flexibility, where multiple logical networks can be provisioned over one common physical infrastructure. This ultimately will lead to lower capex because of the reduction in device complexity and number of devices. Similarly, it will open lower because the operations team will have fewer devices to manage.
- **Support a simplified and flexible integrated security:** Network virtualization also promotes flexible security designs by allowing the use of separate security policies per logical or virtualized entity, where users’ groups and services can be logically separated.
- **Design and operational simplicity:** Network virtualization simplifies the design and provision of path and traffic isolation per application, group, service, and various other logical instances that require end-to-end path isolation.

Note It is important that network designers understand the drivers toward adopting network virtualization from a business point of view, along with the strengths and weaknesses of each design model. This ensures that when a network virtualization concept is considered on a given area within the network or across the entire network, it will deliver the promised value (and not to be used only because it is easy to implement or it is an innovative approach). As discussed earlier in this book, the design that does not address the business’s functional requirements is considered a poor design; consider the design principle “no gold plating” discussed in Chapter 1, “Network Design Requirements: Analysis and Design Principles.”

Note One of the main concerns about network virtualization is the concept of fate sharing, because any failure in the physical network can lead to a failure of multiple virtual networks running over the same physical infrastructure. Therefore, when the network virtualization concept is used, ideally a reliable and highly available network design should be considered as well. Besides the concerns about virtual network availability, there is always a concern about network virtualization (multitenant environment) where multiple virtual networks (VNs) operate over a single physical

network infrastructure and each VN probably has different traffic requirements (different applications and utilization patterns). Therefore, there is a higher potential of having traffic congestion and degraded application quality and user experience if there is no efficient planning with regard to the available bandwidth, number of VNs, traffic volume per VN, applications in use, and the characteristics of the applications. In other words, if there is no adequate bandwidth available and quality of service (QoS) policies to optimize and control traffic behaviors, one VN may overutilize the available bandwidth of the underlying physical network infrastructure. This will usually lead to traffic congestion because other VNs are using the same underlying physical network infrastructure, resulting in what is known as *fate sharing*.

This section covers the primary network virtualization technologies and techniques that you can use to serve different requirements by highlighting the pros and cons of each technology and design approach. This can help network designers (CCDE candidates) to select the best suitable design after identifying and evaluating the different design requirements (business and functional requirements). This section primarily focuses on network virtualization over the enterprise campus network. Chapter 4, “Enterprise Edge Architecture Design,” expands on this topic to cover network virtualization design options and considerations over the WAN.

Network Virtualization Design Elements

As illustrated in Figure 3-9, the main elements in an end-to-end network virtualization design are as follows:

- **Edge control:** This element represents the network access point. Typically, it is a host or end-user access (wired, wireless, or virtual private network [VPN]) to the network where the identification (authentication) for physical to logical network mapping can occur. For example, a contracting employee might be assigned to VLAN X, whereas internal staff is assigned to VLAN Y.
- **Transport virtualization:** This element represents the transport path that will carry different virtualized networks over one common physical infrastructure, such as an overlay technology like a generic routing encapsulation (GRE) tunnel. The terms *path isolation* and *path separation* are commonly used to refer to transport virtualization. Therefore, these terms are used interchangeably throughout this book.
- **Services virtualization:** This element represents the extension of the network virtualization concept to the services edge, which can be shared services among different logically isolated groups, such as an Internet link or a file server located in the data center that must be accessed by only one logical group (business unit).

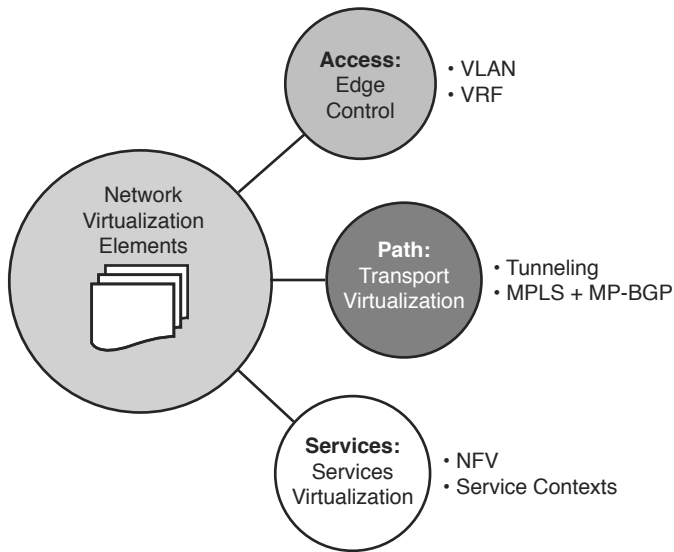


Figure 3-9 *Network Virtualization Elements*

Enterprise Network Virtualization Deployment Models

Now that you know the different elements that, individually or collectively, can be considered as the foundational elements to create network virtualization within the enterprise network architecture, this section covers how you can use these elements with different design techniques and approaches to deploy network virtualization across the enterprise campus. This section also compares these different design techniques and approaches.

Network virtualization can be categorized into the following three primary models, each of which has different techniques that can serve different requirements:

- Device virtualization
- Path isolation
- Services virtualization

Moreover, you can use the techniques of the different models individually to serve certain requirements or combined together to achieve one cohesive end-to-end network virtualization solution. Therefore, network designers must have a good understanding of the different techniques and approaches, along with their attributes, to select the most suitable virtualization technologies and design approach for delivering value to the business.

Device Virtualization

Also known as *device partitioning*, device virtualization represents the ability to virtualize the data plane, control plane, or both, in a certain network node, such as a switch or a router. Using device level virtualization by itself will help to achieve separation at Layer 2, Layer 3, or both, on a local device level. The following are the primary techniques used to achieve device level network virtualization:

- **Virtual LAN (VLAN):** VLAN is the most common Layer 2 network virtualization technique. It is used in every network where one single switch can be divided into multiple logical Layer 2 broadcast domains that are virtually separated from other VLANs. You can use VLANs at the network edge to place an endpoint into a certain virtual network. Each VLAN has its own MAC forwarding table and spanning-tree instance (Per-VLAN Spanning Tree [PVST]).
- **Virtual routing and forwarding (VRF):** VRFs are conceptually similar to VLANs, but from a control plane and forwarding perspective on a Layer 3 device. VRFs can be combined with VLANs to provide a virtualized Layer 3 gateway service per VLAN. As illustrated in Figure 3-10, each VLAN over a 802.1Q trunk can be mapped to a different subinterface that is assigned to a unique VRF, where each VRF maintains its own forwarding and routing instance and potentially leverages different VRF-aware routing protocols (for example, OSPF or EIGRP instance per VRF).




Figure 3-10 *Virtual Routing and Forwarding*

Path Isolation

Path isolation refers to the concept of maintaining end-to-end logical path transport separation across the network. The end-to-end path separation can be achieved using the following main design approaches:

- **Hop by hop:** This design approach, as illustrated in Figure 3-11, is based on deploying end-to-end (VLANs + 802.1Q trunk links + VRFs) per device in the traffic path. This design approach offers a simple and reliable path separation solution. However, for large-scale dynamic networks (large number of virtualized networks), it will be a complicated solution to manage. This complexity is associated with design scalability limitation.




Figure 3-11 *Hop-by-Hop Path Virtualization*

- **Multihop:** This approach is based on using tunneling and other overlay technologies to provide end-to-end path isolation and carry the virtualized traffic across the network. The most common proven methods include the following:
 - **Tunneling:** Tunneling, such as GRE or multipoint GRE (mGRE) (dynamic multipoint VPN [DMVPN]), will eliminate the reliance on deploying end-to-end VRFs and 802.1Q trunks across the enterprise network, because the virtualized traffic will be carried over the tunnel. This method offers a higher level of scalability as compared to the previous option and with simpler operation to some extent. This design is ideally suitable for scenarios where only a part of the network needs to have path isolation across the network.

However, for large-scale networks with multiple logical groups or business units to be separated across the enterprise, the tunneling approach can add complexity to the design and operations. For example, if the design requires path isolation for a group of users across two “distribution blocks,” tunneling can be a good fit, combined with VRFs. However, mGRE can provide the same transport and path isolation goal for larger networks with lower design and operational complexities. (See the section “WAN Virtualization,” in Chapter 4 for a detailed comparison between the different path separation approaches over different types of tunneling mechanisms.)

- **MPLS VPN:** By converting the enterprise to be like a service provider type of network, where the core is Multiprotocol Label Switching (MPLS) enabled and the distribution layer switches to act as provider edge (PE) devices. As in service provider networks, each PE (distribution block) will exchange VPN routing over MP-BGP sessions, as shown in Figure 3-12. (The route reflector [RR] concept can be introduced, as well, to reduce the complexity of full-mesh MP-BGP peering sessions.)

Furthermore, L2VPN capabilities can be introduced in this architecture, such as Ethernet over MPLS (EoMPLS), to provide extended Layer 2 communications across different distribution blocks if required. With this design approach, the end-to-end virtualization and traffic separation can be simplified to a very large extent with a high degree of scalability. (All the MPLS design considerations

and concepts covered in the Service Provider part—Chapter 5, “Service Provider Network Architecture Design,” and Chapter 6, “Service Provider MPLS VPN Services Design,” —in this book are applicable if this design model is adopted by the enterprise.)

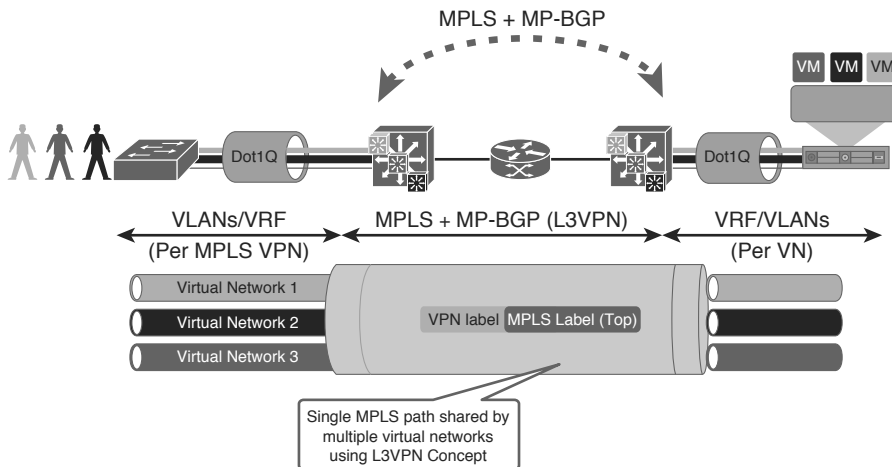


Figure 3-12 *MPLS VPN-Based Path Virtualization*

Figure 3-13 illustrates a summary of the different enterprise campus network’s virtualization design techniques.




Figure 3-13 *Enterprise Campus Network Virtualization Techniques*

As mentioned earlier in this section, it is important for network designers to understand the differences between the various network virtualization techniques. Table 3-3 compares these different techniques in a summarized way from different design angles.

Table 3-3 *Network Virtualization Techniques Comparison*

	End to End (VLAN + 802.1Q + VRF)	VLANs + VRFs + GRE Tunnels	VLANs + VRFs + mGRE Tunnels	MPLS Core with MP-BGP
Scalability	Low	Low	Moderate	High
Operational complexity	High	Moderate	Moderate	Moderate to high
Design flexibility	Low	Moderate	Moderate	High
Architecture	Per hop end-to-end virtualization	P2P (multihop end-to-end virtualization)	P2MP (multihop end-to-end virtualization)	MPLS-L3VPN-based virtualization
Operation staff routing expertise	Basic	Medium	Medium	Advanced
Ideal for	Limited NV scope in terms of size and complexity	Interconnecting specific blocks with NV or as an interim solution	Medium to large overlaid NV design	Large to very large (global scale) end-to-end NV design

Service Virtualization

One of the main goals of virtualization is to separate services access into different logical groups, such as user groups or departments. However, in some scenarios, there may be a mix of these services in term of service access, in which some of these services must only be accessed by a certain group and others are to be shared among different groups, such as a file server in the data center or Internet access, as shown in Figure 3-14.

Therefore, in scenarios like this where service access has to be separated per virtual network or group, the concept of network virtualization must be extended to the services access edge, such as a server with multiple VMs or an Internet edge router with single or multiple Internet links.

Note The virtualization of a network can be extended to other network service appliances, such as firewalls. For instance, you can have a separate virtual firewall per virtual network, to facilitate access control between the virtual user network and the virtualized services and workload, as shown in Figure 3-15. The virtualization of network services appliance can be considered as a “one-to-many” network device level virtualization.

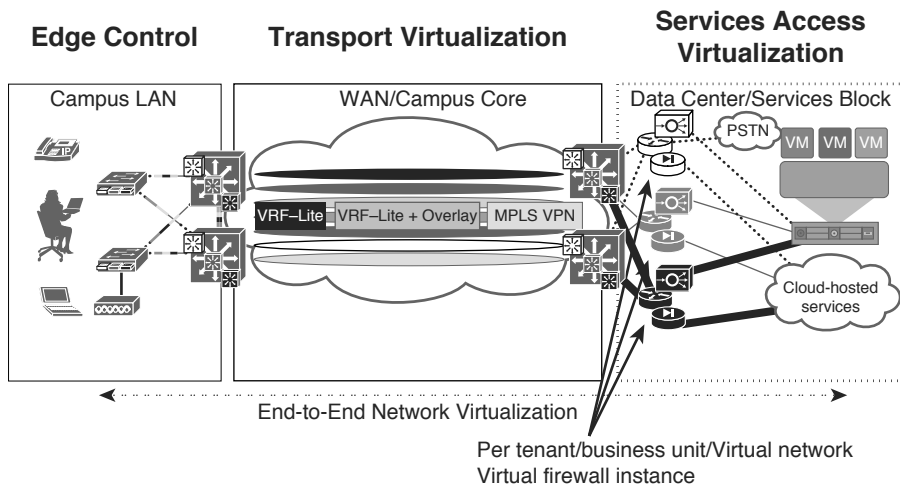


Figure 3-14 *End-to-end Path and Services Virtualization*

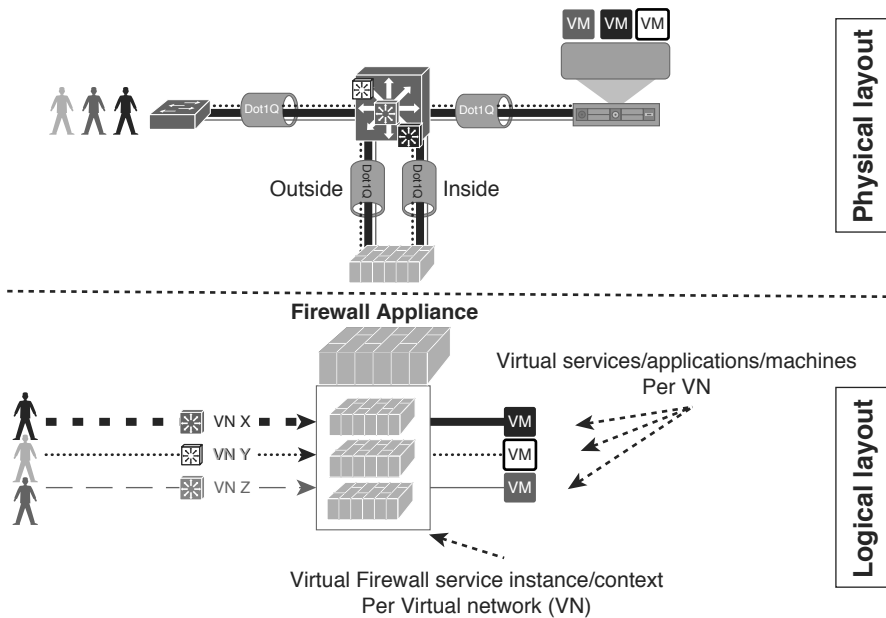


Figure 3-15 *Firewall Virtual Instances*

Furthermore, in multitenant network environments, multiple security contexts offer a flexible and cost-effective solution for enterprises (and for service providers). This approach enables network operators to partition a single pair of redundant firewalls or a single firewall cluster into multiple virtual firewall instances per business unit or tenant. Each tenant can then deploy and manage its own security policies and service access, which are virtually separated. This approach also allows controlled intertenant

communication. For example, in a typical multitenant enterprise campus network environment with MPLS VPN (L3VPN) enabled at the core, traffic between different tenants (VPNs) is normally routed via a firewalling service for security and control (who can access what), as illustrated in Figure 3-16.

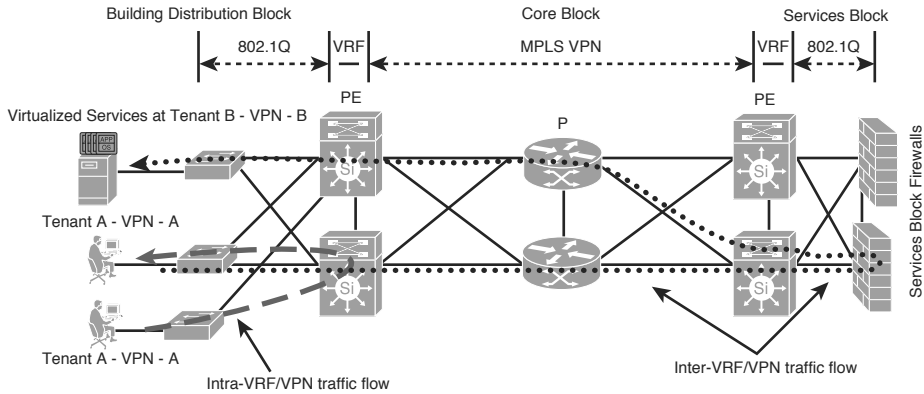


Figure 3-16 Intertenant Services Access Traffic Flow

Figure 3-17 zooms in on the firewall services contexts to show a more detailed view (logical/virtualized view) of the traffic flow between the different tenants/VPNs (A and B), where each tenant has its own virtual firewall service instance located at the services block (or at the data center) of the enterprise campus network.

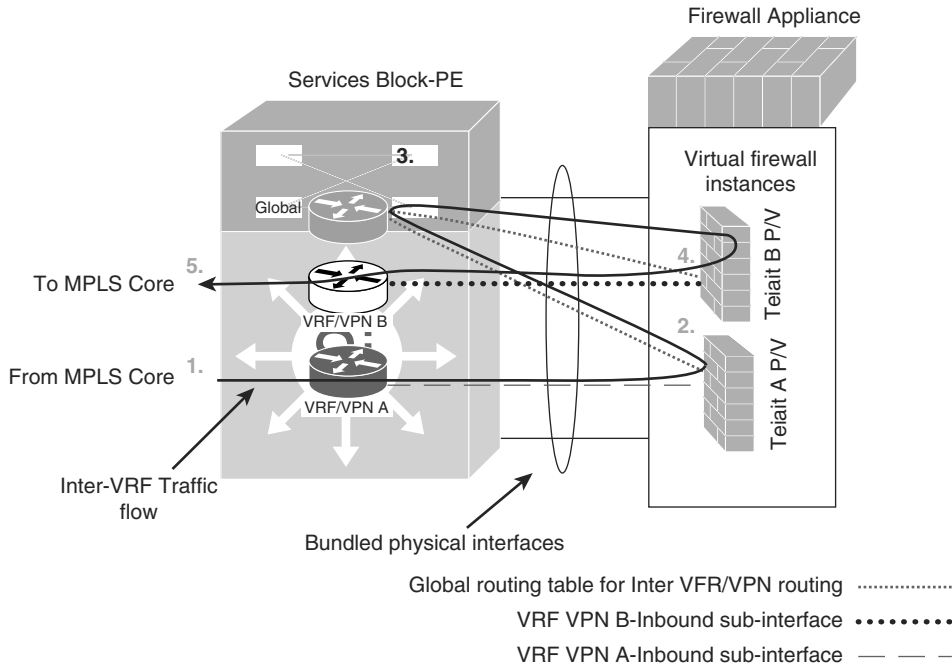


Figure 3-17 Intertenant Services Access Traffic Flow with Virtual Firewall Instances

In addition, the following are the common techniques that facilitate accessing shared applications and network services in multitenant environments:

- VRF-Aware Network Address Translation (NAT):** One of the common requirements in today's multitenant environments with network and service virtualization enabled, is to provide each virtual (tenant) network the ability to access certain services (shared services) either hosted on premise (such as at the enterprise data center or services block) or hosted externally (in a public cloud). Also, providing Internet access to the different tenants (virtual) networks, is a common example of today's multitenant network requirements. To maintain traffic separation between the different tenants (virtual networks) where private IP address overlapping is a common attribute in this type of environment, NAT is considered one of the common and cost-effective solutions to provide NAT per tenant without compromising path separation requirements between the different tenants' networks (virtual networks). When NAT is combined with different virtual network instances (VRFs), it is commonly referred to as *VRF-Aware NAT*, as shown in Figure 3-18.

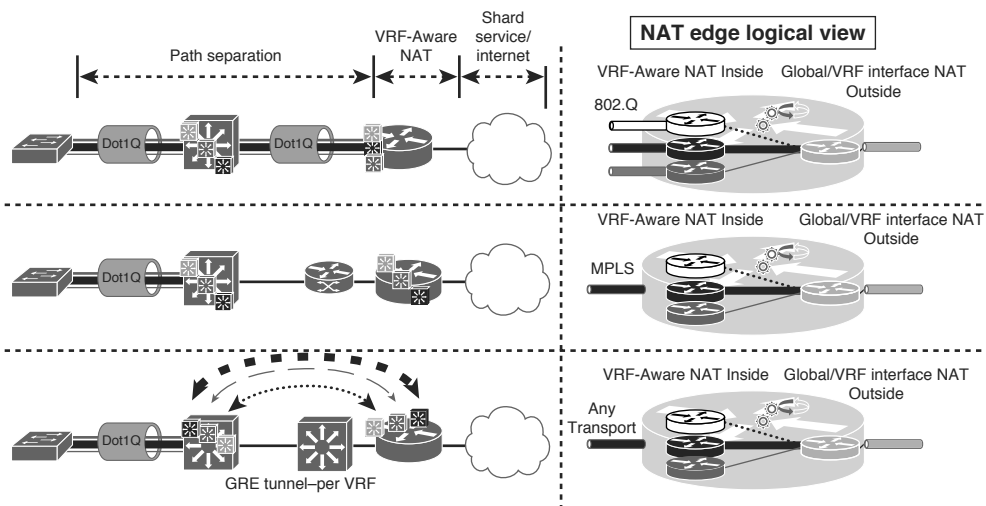


Figure 3-18 VRF-Aware NAT

Note *VRF-aware service infrastructure (VASI)* refers to the ability of an infrastructure or a network node, such as a router, to facilitate the application of features and management services (such as encryption and NAT) between VRFs internally within the same node, using virtual interfaces. For two VRFs to communicate internally within a network node (router), a VASI virtual interface pair can be configured. Each interface in this pair must be associated with a different VRF so that those two virtual interfaces can be logically wired,¹ as illustrated in Figure 3-19. This capability is available in some high-end Cisco routers.

1. "Configuring VRF-Aware Service Infrastructure," http://www.cisco.com/c/en/us/td/docs/ios_xr_sw/iosxr_r3-8/vfw/configuration/guide/vfc38/vfc38vas.pdf

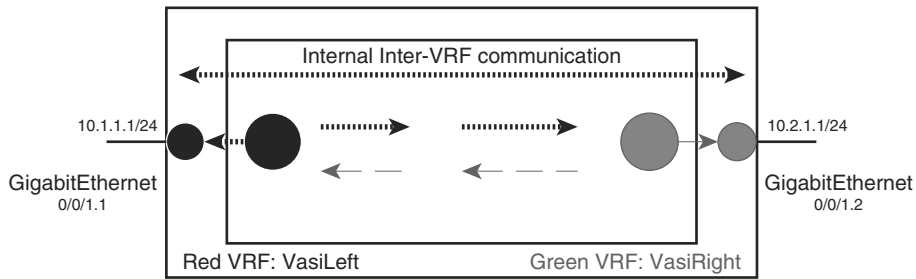


Figure 3-19 VRF-aware Services Infrastructure

- **Network function virtualization (NFV):** The concept of NFV is based on virtualizing network functions that typically require a dedicated physical node, appliances, or interfaces. In other words, NFV can potentially take any network function typically residing in purpose-built hardware and abstract it from that hardware. As depicted in Figure 3-20, this concept offers businesses several benefits, including the following:
 - Reduce the total cost of ownership (TCO) by reducing the required number and diversity of specialized appliances
 - Reduce operational cost (for example, less power and space)
 - Offer a cost-effective capital investment
 - Reduce the level of complexity of integration and network operations
 - Reduce time to market for the business by offering the ability to enable specialized network services (Especially in multitenant where a separate network function/service per tenant can be provisioned faster)




Figure 3-20 NFV Benefits

This concept helps businesses to adopt and deploy new services quickly (faster time to market), and is consequently considered a business innovation enabler. This is simply

because purpose-built hardware functionalities have now been virtualized, and it is a matter of service enablement rather than relying on new hardware (along with infrastructure integration complexities).

Note The concept of NFV is commonly adopted by service provider networks nowadays. Nonetheless, this concept is applicable and usable in enterprise networks and enterprise data center networks that want to gain its benefits and flexibility.

Note In large-scale networks with a very high volume of traffic (typically carrier grade), hardware resource utilization and limits must be considered.

Summary

The enterprise campus is one of the vital parts of the modular enterprise network. It is the medium that connects the end users and the different types of endpoints such as printers, video endpoints, and wireless access points to the enterprise network. Therefore, having the right structure and design layout that meets current and future requirements is critical, including the physical infrastructure layout, Layer 2, and Layer 3 designs. To achieve a scalable and flexible campus design, you should ideally base it on hierarchical and modular design principles that optimize the overall design architecture in terms of fault isolation, simplicity, and network convergence time. It should also offer a desirable level of flexibility to integrate other networks and new services and to grow in size.

However, the concept of network virtualization helps enterprises to utilize the same underlying physical infrastructure while maintaining access, and path and services access isolation, to meet certain business goals or functional security requirements. As a result, enterprises can lower capex and opex and reduce the time and effort required to provision a new service or a new logical network. However, the network designer must consider the different network virtualization design options, along with the strengths and weaknesses of each, to deploy the suitable network virtualization technique that meets current and future needs. These needs must take into account the different variables and constraints, such as staff knowledge and the hardware platform supported features and capabilities.

Further Reading

“Borderless Campus 1.0 Design Guide,” <http://www.cisco.com>

“Enterprise Campus 3.0 Architecture: Overview and Framework,” <http://www.cisco.com>

“Cisco Wired LAN, Cisco Validated Design,” <http://www.cisco.com>

“Network Virtualization Solutions, Design Guides,” <http://www.cisco.com>

“Network Services Virtualization,” <http://www.cisco.com>

This page intentionally left blank