



Data Center Virtualization Fundamentals



ciscopress.com

Gustavo Santana

FREE SAMPLE CHAPTER



SHARE WITH OTHERS

Data Center Virtualization Fundamentals

Gustavo Alessandro Andrade Santana, CCIE No. 8806

Cisco Press

800 East 96th Street

Indianapolis, IN 46240

Data Center Virtualization Fundamentals

Copyright © 2014 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.

Library of Congress Control Number: 2013940880

Printed in the United States of America

First Printing June 2013

ISBN-13: 978-1-58714-324-3

ISBN-10: 1-58714-324-0

Warning and Disclaimer

This book is designed to provide information about data center technologies. Every effort has been made to make this book as complete and as accurate as possible, but no warranty or fitness is implied.

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

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

Trademark Acknowledgments

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

Corporate and Government Sales

The publisher offers excellent discounts on this book when ordered in quantity for bulk purchases or special sales, which may include electronic versions and/or custom covers and content particular to your business, training goals, marketing focus, and branding interests. For more information, please contact: **U.S. Corporate and Government Sales** 1-800-382-3419 corpsales@pearsontechgroup.com

For sales outside of the U.S. please contact: **International Sales** international@pearsoned.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 e-mail at feedback@ciscopress.com. Please make sure to include the book title and ISBN in your message.

We greatly appreciate your assistance.

Publisher: Paul Boger

Business Operation Manager, Cisco Press: Jan Cornelissen

Associate Publisher: Dave Dusthimer

Executive Editor: Mary Beth Ray

Development Editor: Eleanor C. Bru

Managing Editor: Sandra Schroeder

Copy Editor: John Edwards

Project Editor: Seth Kerney

Technical Editors: Maurilio Gorito, Krishna Arji

Editorial Assistant: Vanessa Evans

Proofreader: Sheri Cain

Cover Designer: Mark Shirar

Indexer: Larry Sweazy

Composition: Jake McFarland



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, CCFR, CCNA, CGNP, CCSP, CCVP, Cisco, the 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, PowerPanel, 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

Gustavo A. A. Santana, CCIE No. 8806, is a Cisco Technical Solutions Architect working in enterprise and service provider data center projects that require a greater integration among multiple technology areas such as networking, application optimization, storage, and servers.

With more than 15 years of experience in the data center industry, Gustavo has led and coordinated a team of specialized Cisco engineers in Brazil. A true believer of education as a technology catalyst, he has also dedicated himself to the technical development of many IT professionals from customer, partner, and strategic alliance organizations.

In addition to holding two CCIE certifications (Routing & Switching and Storage Networking), Gustavo is also a VMware Certified Professional (VCP) and an SNIA Certified Storage Networking Expert (SCSN-E). A frequent speaker at Cisco and data center industry events, he holds a degree in computer engineering from Instituto Tecnológico de Aeronáutica (ITA-Brazil) and an MBA in strategic IT management from Fundação Getúlio Vargas (FGV-Brazil).

Gustavo maintains a personal blog in which he discusses topics related to data center virtualization technologies at <http://gustavoasantana.net>.

About the Technical Reviewers

Maurilio Gorito, CCIE, is the certification manager at Riverbed Technology Institute, leading Riverbed Technology's Certification Program. Previously, he managed the Cisco CCIE Routing and Switching certification program, and he is a triple Cisco Certified Internetwork Expert (CCIE). Maurilio has 25 years of experience in the IT field and more than 20 years of combined experience in education, testing, and assessment. Maurilio has written a book, written articles, and reviewed several technical books for Cisco Press. Maurilio holds bachelor's degrees in education, pedagogy, and mathematics from Centro Universitario Geraldo Di Biasi, Brazil. Maurilio is currently serving on the board of directors at Performance Testing Council (PTC).

Krishna Arji is a senior manager at Cisco. In this role, he is responsible for the development of technology that enables the delivery of Cisco services. Krishna has held various positions in the Cisco Services Technology Group at Cisco. His current assignment is to enable delivery for the Cisco BYOD service offerings. In the past, he played a key role in evaluating and developing technologies required for the delivery of cloud planning, design, and implementation services. Under his leadership, his team developed several technologies to perform routing, switching, data center, security, and WLAN assessments of customers' infrastructures. His areas of expertise include networking, software design and development, and data center technologies such as virtualization. Krishna holds a bachelor's degree in electronics and communications engineering, and he has a master's degree in enterprise software technologies. He is currently enrolled in an MBA program at Haas School of Business, University of California, Berkeley. He has a patent pending with USPTO for Automated Assessments of Storage Area Networks (Serial No. 13/115, 141).

Dedications

This book is dedicated to my wife and true love, Carlene, whose sacrifice and unconditional support were crucial to this endeavor, and to my lovely daughter, Carolina, whose one-year-old curiosity constantly inspired me to go one step further.

I also dedicate this book to my parents, Honorio and Cleia, who have taught me that one can only learn by being fearless and humble.

Finally, this book is also dedicated to every person who has devoted efforts to the rewarding experience of teaching someone.

Acknowledgments

The process of creating a book can be aptly defined as a “sponsored solitude.” Certainly, the uncountable lonely writing hours would be fruitless without the support of an entire network of relatives, friends, and professionals who are acknowledged here.

First, I would like to thank my sister Raquel and brother André for the family support during this book writing.

I would also like to express my gratitude to my friend and trusted advisor Alexandre M. S. P. Moraes, who has shared invaluable opinions and insights since the very early stages of this book.

Many thanks to Andrey Lee for the wonderful illustrations in Chapters 1 and 17.

Sincere thanks to Paulo Quinta, Fernanda Spinardi, and Marcelo Ehalt for helping me coordinate my professional life and this writing.

My thanks to the technical reviewers Maurilio Gorito and Krishna Arji for their active contributions and focus to make this work more effective for its targeted readership.

A personal thanks to the Brazilian data center tiger team, which has always served as my favorite “think tank” for best practices and the exchange of experiences.

I would also like to thank two very talented instructors from Firefly: Dan Murray and Fabricio Grimaldi.

I am also very grateful to the people who have contributed with the equipment used in this publication: Shane Hudson and Bilal El-Ayi (from GoldLabs), Ohad Richberg (and his amazing CPOC Israel team), François Tallet and Mark Allen (from the Cisco Enterprise Core Business Unit), and Hugo Marques.

Thanks to all the Pearson production team, especially Ellie Bru and Seth Kerney, who helped me to create the final version of this book.

A special thank-you goes to Mary Beth Ray and Anand Sundaram for supporting the idea of a data center book with a different approach.

Contents at a Glance

Foreword xxiii

Introduction xxv

Part I What Is Virtualization?

Chapter 1 Virtualization History and Definitions 1

Part II Virtualization in Network Technologies

Chapter 2 Data Center Network Evolution 25

Chapter 3 The Humble Beginnings of Network Virtualization 45

Chapter 4 An Army of One: ACE Virtual Contexts 109

Chapter 5 Instant Switches: Virtual Device Contexts 183

Chapter 6 Fooling Spanning Tree 231

Chapter 7 Virtualized Chassis with Fabric Extenders 287

Chapter 8 A Tale of Two Data Centers 319

Part III Virtualization in Storage Technologies

Chapter 9 Storage Evolution 387

Chapter 10 Islands in the SAN 409

Chapter 11 Secret Identities 453

Chapter 12 One Cable to Unite Us All 493

Part IV Virtualization in Server Technologies

Chapter 13 Server Evolution 559

Chapter 14 Changing Personalities 581

Chapter 15 Transcending the Rack 657

Chapter 16 Moving Targets 735

Part V End-to-End Virtualization

Chapter 17 The Virtual Data Center and Cloud Computing 785

Part VI Appendixes

Appendix A Cisco Data Center Portfolio 809

Appendix B IOS, NX-OS, and Application Control Software Command-Line Interface Basics 847

Index 873

Contents

Foreword xxiii

Introduction xxv

Part I **What Is Virtualization?**

Chapter 1 **Virtualization History and Definitions** **1**

Data Center Essential Definitions 2

 Data Center Evolution 3

 Operational Areas and Data Center Architecture 5

The Origins of Data Center Virtualization 8

 Virtual Memory 8

 Mainframe Virtualization 10

 Hot Standby Router Protocol 11

 Defining Virtualization 12

 Data Center Virtualization Timeline 12

Classifying Virtualization Technologies 14

 A Virtualization Taxonomy 15

 Virtualization Scalability 17

 Technology Areas 18

 Classification Examples 21

Summary 22

Further Reading 22

Part II **Virtualization in Network Technologies**

Chapter 2 **Data Center Network Evolution** **25**

Ethernet Protocol: Then and Now 26

 Ethernet Media 27

 Coaxial Cable 27

 Twisted-Pair 28

 Optical Fiber 30

 Direct-Attach Twinaxial Cables 32

 Ethernet Data Rate Timeline 33

Data Center Network Topologies 34

 Data Center Network Layers 35

 Design Factors for Data Center Networks 36

 Physical Network Layout Considerations 39

 The ANSI/TIA-942 Standard 40

| | |
|---|----|
| Network Virtualization Benefits | 42 |
| Network Logical Partitioning | 42 |
| Network Simplification and Traffic Load Balancing | 43 |
| Management Consolidation and Cabling Optimization | 44 |
| Network Extension | 44 |
| Summary | 44 |
| Further Reading | 44 |

Chapter 3 The Humble Beginnings of Network Virtualization 45

| | |
|---|-----|
| Network Partitioning | 47 |
| Concepts from the Bridging World | 47 |
| Defining VLANs | 49 |
| VLAN Trunks | 52 |
| Two Common Misconceptions About VLANs | 56 |
| Misconception Number 1: A VLAN Must Be Associated to an IP Subnet | 56 |
| Misconception Number 2: Layer 3 VLANs | 58 |
| Spanning Tree Protocol and VLANs | 61 |
| Spanning Tree Protocol at Work | 63 |
| Port States | 70 |
| Spanning Tree Protocol Enhancements | 72 |
| Spanning Tree Instances | 74 |
| Private VLANs | 78 |
| VLAN Specifics | 83 |
| Native VLAN | 84 |
| Reserved VLANs IDs | 84 |
| Resource Sharing | 85 |
| Control and Management Plane | 85 |
| Concepts from the Routing World | 87 |
| Overlapping Addresses in a Data Center | 87 |
| Defining and Configuring VRFs | 90 |
| VRFs and Routing Protocols | 92 |
| VRFs and the Management Plane | 98 |
| VRF-Awareness | 100 |
| VRF Resource Allocation Control | 101 |
| Use Case: Data Center Network Segmentation | 103 |
| Summary | 105 |
| Further Reading | 107 |

| | | |
|------------------|--|------------|
| Chapter 4 | An Army of One: ACE Virtual Contexts | 109 |
| | Application Networking Services | 111 |
| | The Use of Load Balancers | 111 |
| | Load-Balancing Concepts | 115 |
| | Layer 4 Switching Versus Layer 7 Switching | 120 |
| | Connection Management | 122 |
| | Address Translation and Load Balancing | 124 |
| | <i>Server NAT</i> | 124 |
| | <i>Dual NAT</i> | 125 |
| | <i>Port Redirection</i> | 126 |
| | <i>Transparent Mode</i> | 126 |
| | Other Load-Balancing Applications | 127 |
| | <i>Firewall Load Balancing</i> | 127 |
| | <i>Reverse Proxy Load Balancing</i> | 128 |
| | Offloading Servers | 130 |
| | <i>SSL Offload</i> | 130 |
| | <i>TCP Offload</i> | 133 |
| | <i>HTTP Compression</i> | 134 |
| | Load Balancer Proliferation in the Data Center | 135 |
| | Load Balancer Performance | 135 |
| | Security Policies | 136 |
| | Suboptimal Traffic | 137 |
| | Application Environment Independency | 138 |
| | ACE Virtual Contexts | 139 |
| | Application Control Engine Physical Connections | 141 |
| | <i>Connecting an ACE Appliance</i> | 141 |
| | <i>Connecting an ACE Module</i> | 144 |
| | Creating and Allocating Resources to Virtual Contexts | 145 |
| | Integrating ACE Virtual Contexts to the Data Center Network | 156 |
| | <i>Routed Design</i> | 156 |
| | <i>Bridged Design</i> | 158 |
| | <i>One-Armed Design</i> | 160 |
| | Managing and Configuring ACE Virtual Contexts | 162 |
| | <i>Allowing Management Traffic to a Virtual Context</i> | 162 |
| | <i>Allowing Load Balancing Traffic Through a Virtual Context</i> | 163 |
| | Controlling Management Access to Virtual Contexts | 171 |

ACE Virtual Context Additional Characteristics 176

Sharing VLANs Among Contexts 177

Virtual Context Fault Tolerance 177

Use Case: Multitenant Data Center 179

Summary 181

Further Reading 182

Chapter 5 Instant Switches: Virtual Device Contexts 183

Extending Device Virtualization 184

Why Use VDCs? 187

VDCs in Detail 188

Creating and Configuring VDCs 190

VDC Names and CLI Prompts 198

Virtualization Nesting 199

Allocating Resources to VDCs 202

Using Resource Templates 211

Managing VDCs 214

VDC Operations 214

Processes Failures and VDCs 216

VDC Out-of-Band Management 217

Role-Based Access Control and VDCs 222

Global Resources 225

Use Case: Data Center Security Zones 225

Summary 227

Further Reading 229

Chapter 6 Fooling Spanning Tree 231

Spanning Tree Protocol and Link Utilization 232

Link Aggregation 234

Server Connectivity and NIC Teaming 238

Cross-Switch PortChannels 240

Virtual PortChannels 241

Virtual PortChannel Definitions 242

Configuring Virtual PortChannels 247

Step 1: Defining the Domain 248

Step 2: Establishing Peer Keepalive Connectivity 248

Step 3: Creating the Peer Link 250

| | |
|--|-----|
| <i>Step 4: Creating the Virtual PortChannel</i> | 252 |
| Spanning Tree Protocol and Virtual PortChannels | 254 |
| <i>Peer Link Failure and Orphan Ports</i> | 258 |
| First-Hop Routing Protocols and Virtual PortChannels | 259 |
| Layer 2 Multipathing and vPC+ | 265 |
| FabricPath Data Plane | 266 |
| FabricPath Control Plane | 269 |
| FabricPath and Spanning Tree Protocol | 272 |
| Virtual PortChannel Plus | 276 |
| Use Case: Evolution of Network PODs | 281 |
| Summary | 285 |
| Further Reading | 286 |

Chapter 7 Virtualized Chassis with Fabric Extenders 287

| | |
|---|-----|
| Server Access Models | 288 |
| Understanding Fabric Extenders | 291 |
| Fabric Extender Options | 295 |
| Connecting a Fabric Extender to a Parent Switch | 296 |
| Fabric Extended Interfaces and Spanning Tree Protocol | 299 |
| Fabric Interfaces Redundancy | 301 |
| Fabric Extender Topologies | 305 |
| Straight-Through Topologies | 305 |
| Dual-Homed Topologies | 309 |
| Use Case: Mixed Access Data Center | 315 |
| Summary | 317 |
| Further Reading | 318 |

Chapter 8 A Tale of Two Data Centers 319

| | |
|--|-----|
| A Brief History of Distributed Data Centers | 321 |
| The Cold Age (Mid-1970s to 1980s) | 321 |
| The Hot Age (1990s to Mid-2000s) | 322 |
| The Active-Active Age (Mid-2000s to Today) | 324 |
| The Case for Layer 2 Extensions | 324 |
| Challenges of Layer 2 Extensions | 325 |
| Ethernet Extensions over Optical Connections | 327 |
| Virtual PortChannels | 328 |
| FabricPath | 330 |

| | |
|---|-----|
| Ethernet Extensions over MPLS | 332 |
| MPLS Basic Concepts | 333 |
| Ethernet over MPLS | 338 |
| Virtual Private LAN Service | 342 |
| Ethernet Extensions over IP | 352 |
| MPLS over GRE | 352 |
| Overlay Transport Virtualization | 354 |
| OTV Terminology | 357 |
| OTV Basic Configuration | 359 |
| OTV Loop Avoidance and Multiboming | 365 |
| Migration to OTV | 366 |
| OTV Site Designs | 373 |
| VLAN Identifiers and Layer 2 Extensions | 377 |
| Internal Routing in Connected Data Centers | 380 |
| Use Case: Active-Active Greenfield Data Centers | 382 |
| Summary | 384 |
| Further Reading | 386 |

Part III Virtualization in Storage Technologies

Chapter 9 Storage Evolution 387

| | |
|---|-----|
| Data Center Storage Devices | 387 |
| Hard Disk Drives | 388 |
| Disk Arrays | 389 |
| Tape Drives and Libraries | 390 |
| Accessing Data in Rest | 391 |
| Block-Based Access | 392 |
| <i>Small Computer Systems Interface</i> | 392 |
| <i>Mainframe Storage Access</i> | 396 |
| <i>Advanced Technology Attachment</i> | 397 |
| File Access | 397 |
| <i>Network File System</i> | 398 |
| <i>Common Internet File System</i> | 398 |
| Record Access | 398 |
| Storage Virtualization | 399 |
| Virtualizing Storage Devices | 402 |
| Virtualizing LUNs | 404 |

| | |
|---------------------------|-----|
| Virtualizing File Systems | 406 |
| Virtualizing SANs | 407 |
| Summary | 408 |
| Further Reading | 408 |

Chapter 10 Islands in the SAN 409

| | |
|---|-----|
| Some Fibre Channel Definitions | 410 |
| Fibre Channel Layers | 411 |
| Fibre Channel Topologies and Port Types | 412 |
| Fibre Channel Addressing | 413 |
| Frames, Sequences, and Exchanges | 415 |
| Flow Control | 417 |
| Classes of Service | 420 |
| Fabric Processes | 420 |
| Fabric Initialization | 422 |
| Fabric Shortest Path First | 424 |
| Register State Change Notification | 426 |
| Fibre Channel Logins | 427 |
| Zoning | 429 |
| Defining and Exploring VSANs | 430 |
| SAN Islands | 430 |
| VSAN Creation | 432 |
| VSAN Trunking | 434 |
| Zoning and VSANs | 439 |
| FSPF and VSANs | 442 |
| VSAN Scoping | 445 |
| Use Case: SAN Consolidation | 447 |
| Summary | 450 |
| Further Reading | 451 |

Chapter 11 Secret Identities 453

| | |
|--|-----|
| Fibre Channel over IP | 454 |
| FCIP High Availability | 460 |
| Use Case: SAN Extension with Traffic Engineering | 462 |
| Inter-VSAN Routing | 464 |
| IVR Infrastructure | 465 |
| IVR Zoning | 467 |
| Use Case: Transit VSAN | 472 |

| | |
|--|-----|
| N_Port Virtualization | 473 |
| Configuring N_Port Virtualization | 476 |
| NPV Traffic Management | 482 |
| Deploying Port WWN Virtualization on NPV | 486 |
| Use Case: Blade Server Hosting Data Center | 488 |
| Summary | 490 |
| Further Reading | 491 |

Chapter 12 One Cable to Unite Us All 493

| | |
|--|-----|
| The Case for Data Center Networking Convergence | 495 |
| Data Center Bridging | 497 |
| Priority-Based Flow Control | 498 |
| Enhanced Transmission Selection | 500 |
| Data Center Bridging eXchange Protocol | 501 |
| Congestion Notification | 503 |
| Introducing Fibre Channel over Ethernet | 504 |
| FCoE Elements | 505 |
| FCoE Initialization Protocol | 507 |
| Deploying Unified Server Access | 509 |
| Configuring Unified Server Access on Single-Context Switches | 510 |
| Configuring Unified Server Access with Storage VDCs | 519 |
| Configuring Multihop FCoE | 523 |
| Configuring Virtual Fibre Channel PortChannels | 528 |
| FCoE N_Port Virtualization | 532 |
| Unified Fabric Designs | 535 |
| Server Access Layer Unified Designs | 536 |
| <i>FCoE and Virtual PortChannels</i> | 538 |
| <i>FCoE and Blade Servers</i> | 540 |
| Beyond the Access Layer | 542 |
| <i>Converged Access Model</i> | 542 |
| <i>Converged Aggregation Model</i> | 543 |
| FCoE and SAN Extension | 545 |
| Use Case: LAN and SAN Management Separation | 546 |
| Summary | 556 |
| Further Reading | 557 |

Part IV Virtualization in Server Technologies**Chapter 13 Server Evolution 559**

- Server Architectures 560
 - Mainframes 560
 - RISC Servers 561
 - x86 Servers 562
- x86 Hardware Evolution 562
 - CPU Evolution 564
 - Memory Evolution 566
 - Expansion Bus Evolution 569
 - Physical Format Evolution 571
- Introducing x86 Server Virtualization 572
 - Virtualization Unleashed 574
- Unified Computing 578
- Summary 580
- Further Reading 580

Chapter 14 Changing Personalities 581

- Server Provisioning Challenges 583
 - Server Domain Operations 584
 - Infrastructure Domain Operations 585
- Unified Computing and Service Profiles 586
- Building Service Profiles 588
 - Identifying a Service Profile 594
 - Storage Definitions 595
 - Network Definitions 599
 - Virtual Interface Placement 602
 - Server Boot Order 604
 - Maintenance Policy 606
 - Server Assignment 606
 - Operational Policies 608
 - Configuration* 608
 - External IPMI Management Configuration* 609
 - Management IP Address* 610
 - Additional Policies* 611

| | |
|---|-----|
| Associating a Service Profile to a Server | 612 |
| Installing an Operating System | 620 |
| Verifying Stateless Computing | 625 |
| Using Policies | 626 |
| BIOS Setting Policies | 627 |
| Firmware Policies | 633 |
| Industrializing Server Provisioning | 637 |
| Cloning | 638 |
| Pools | 639 |
| Service Profile Templates | 640 |
| Server Pools | 649 |
| Use Case: Seasonal Workloads | 653 |
| Summary | 655 |
| Further Reading | 656 |

Chapter 15 Transcending the Rack 657

| | |
|--|-----|
| Introduction to Virtual Networking | 658 |
| Virtual Switch Challenges | 660 |
| Cisco Nexus 1000V Architecture | 661 |
| Nexus 1000V Communication Modes | 663 |
| Port Profiles and Dynamic Interface Provisioning | 664 |
| Deploying Nexus 1000V | 666 |
| External Connectivity and Link Aggregation | 684 |
| NX-OS Features in the Virtual World | 688 |
| MAC Address Table | 691 |
| Access Lists | 692 |
| Online Migrations and Nexus 1000V | 693 |
| Virtual Extensible Local Area Networks | 697 |
| Introducing Virtual Machine Fabric Extender | 705 |
| Deploying VM-FEX | 707 |
| Enabling Dynamic vNICs on a UCS Service Profile | 707 |
| Preparing VMware vSphere Host to Deploy VM-FEX | 709 |
| Using the UCS Manager VMware Integration Wizard | 711 |
| Migrating Virtual Machines to VM-FEX | 716 |
| Online Migrations and VM-FEX | 720 |
| VM-FEX High-Performance Mode | 723 |

Use Case: Data Center Merging 731

Summary 733

Further Reading 734

Chapter 16 Moving Targets 735

Virtual Network Services Definitions 736

Virtual Network Services Data Path 738

vPath-Enabled Virtual Network Services 740

Cisco Virtual Security Gateway: Compute Virtual Firewall 742

Installing Virtual Security Gateway 743

Creating Security Policies 745

Sending Data Traffic to VSG 747

Virtual Machine Attributes and Virtual Zones 751

Cisco ASA 1000V: Edge Virtual Firewall 754

Installing ASA 1000V 755

Sending Data Traffic to ASA 1000V 758

Configuring Security Policies on ASA 1000V 761

Application Acceleration 763

WAN Acceleration and Online Migration 769

Routing in the Virtual World 771

Site Selection and Server Virtualization 775

Route Health Injection 775

Global Server Load Balancing 777

Location/ID Separation Protocol 779

Use Case: Virtual Data Center 781

Summary 783

Further Reading 784

Part V End-to-End Virtualization

Chapter 17 The Virtual Data Center and Cloud Computing 785

The Virtual Data Center 786

Automation and Standardization 789

What Is Cloud Computing? 793

Cloud Implementation Example 797

Journey to the Cloud 799

Networking in the Clouds 800

| | |
|--------------------------------|-----|
| Software-Defined Networks | 800 |
| OpenStack | 801 |
| Network Overlays | 802 |
| Cisco Open Network Environment | 804 |
| Before We Go... | 805 |
| Summary | 806 |
| Further Reading | 807 |

Part VI: Appendixes

Appendix A Cisco Data Center Portfolio 809

| | |
|---|-----|
| Cisco Application Control Engine | 809 |
| Cisco Adaptive Security Appliances 5585-X | 811 |
| Cisco ASA 1000V Cloud Firewall | 812 |
| Cisco Catalyst 6500 Series Switches | 813 |
| Cisco Cloud Portal | 816 |
| Cisco Intelligent Automation Solutions | 817 |
| Automation Software Components | 817 |
| Cisco Intelligent Automation for Cloud Solution | 819 |
| Cisco Intelligent Automation for SAP | 820 |
| Cisco MDS 9000 Series Multilayer Switches | 820 |
| Cisco Prime Network Analysis Module | 823 |
| Cisco Nexus Data Center Switches | 823 |
| Cisco Nexus 1000V Series Switches | 824 |
| Nexus 1010 and 1100 Virtual Services Appliances | 824 |
| Cisco Nexus 2000 Series Fabric Extenders | 825 |
| Cisco Nexus 3000 Series Switches | 827 |
| Cisco Nexus 4000 Series Switches | 828 |
| Cisco Nexus 5000 and 5500 Series Switches | 829 |
| Cisco Nexus 6000 Series Switches | 831 |
| Cisco Nexus 7000 Series Switches | 832 |
| Cisco Unified Computing System | 835 |
| Cisco 6100 and 6200 Series Fabric Interconnects | 836 |
| Cisco UCS 5100 Series Blade Server Chassis | 836 |
| Cisco UCS 2100 and 2200 Series Fabric Extenders | 837 |
| Cisco UCS B-Series Blade Servers | 837 |
| Cisco UCS C-Series Rack Servers | 838 |

| | |
|---|-----|
| Cisco UCS Virtual Interface Cards | 839 |
| Unified Management Solutions | 840 |
| Cisco Application Network Manager | 840 |
| Cisco Prime Data Center Network Manager | 841 |
| Cisco UCS Manager and UCS Central | 842 |
| Virtual Network Management Center | 843 |
| Virtual Security Gateway | 843 |
| Virtualization Techniques Mapping | 844 |
| Further Reading | 844 |

Appendix B IOS, NX-OS, and Application Control Software Command-Line Interface Basics 847

| | |
|--|-----|
| IOS Command-Line Interface Basics | 847 |
| Command Modes | 848 |
| Getting Context-Sensitive Help | 850 |
| Abbreviating Commands and Using Shortcuts | 854 |
| Managing Configuration Files | 855 |
| Using Debug Commands | 858 |
| NX-OS Command-Line Interface | 859 |
| NX-OS Access | 860 |
| NX-OS Modularity | 861 |
| NX-OS and Running Configuration Files | 863 |
| NX-OS Command-Line Interface Optimizations | 866 |
| Configuration Version Management, Batches, and Scripts | 866 |
| Application Control Software Command-Line Interface | 870 |

Index 873

Icons Used in This Book



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).
- Italics indicate arguments for which you supply actual values.
- Vertical bars (|) separate alternative, mutually exclusive elements.
- Square brackets [] indicate optional elements.
- Braces { } indicate a required choice.
- Braces within brackets [{ } indicate a required choice within an optional element.
- Introduction

Foreword

With the rapid growth of the Internet economy and the explosion of information technology, the data center is playing a pivotal role and is one of the most exciting fields in the world of IT today. The trend continues with both virtualization and cloud computing fueling growth and making data center solutions more efficient and scalable. More specifically, organizations using virtualization technologies are seeing greater returns and more viability to deal with the growing demands of the economy.

Data center virtualization is an evolutionary process that was started several years ago within mainframe computer rooms, and it has dramatically intensified in the last few years. Its proposed freedom from physical boundaries has produced benefits in each technology area, and much more importantly, from an architectural perspective.

However, due to these environments' increasing complexity, a data center professional must possess a challenging breadth of knowledge in several different areas, such as networking, storage, servers, operating systems, application, and security.

Data Center Virtualization Fundamentals is a comprehensive book that introduces virtualization technologies in data center environments, encompassing all these knowledge areas. It does not take a product-based approach as many others do, but an architectural one, offering theoretical concepts, illustrative configurations, and real-world designs for each virtualization technique. The book provides a first step for students and professionals who want to understand the state of data center technologies today. And in my opinion, virtualization technologies are the best way to achieve this feat because one must be aware of the physical challenges of data center environments before learning such techniques.

There is a lot of misconception when talking about virtualization, and people immediately think of it in the context of virtual servers. However, virtualization is not restricted to a single technology area in the data center. This book intends to make an account of the main data center virtualization technologies, revealing their impact and applicability to these environments as a whole. It encourages readers to escape their technical comfort-zone and learn how each decision may impact other data center teams. A strong knowledge of the theoretical basis of the data center is necessary to walk amidst clouds, and this is exactly what this book brings.

Author Gustavo A. A. Santana is a seasoned expert with years of experience, and has done a superb job putting this material together. He has demonstrated his skills and command of the technology, using a unique approach in translating the most complex and highly technical information into simple, easy-to-understand material. Readers will definitely appreciate this book.

Finally, this book is an essential reference and will be a valuable asset for potential candidates pursuing their Cisco Data Center certifications. I am confident that in reading this book, individuals will inevitably gain extensive knowledge and hands-on experience during their certification preparations. If you're looking for a truly comprehensive guide to virtualization, this is the one!

Yusuf Bhajji

Senior Manager, Expert Certifications (CCIE, CCDE, CCAr)

Learning@Cisco

“If you can't explain it simply, you don't understand it well enough.”

—Albert Einstein

Introduction

“I am very interested in learning data center technologies. How should I start?”

Since I first heard this question, I have seen many IT professionals become overwhelmed with the vertigo-inducing development of new data center technologies. From my perspective, their frustration was mainly caused by attempting to understand this subject without being properly introduced to the most fundamental concepts and definitions related to these complex environments. And that opinion has always formed the basis of my advice to them.

However, as the years passed, I observed how my answer to this question was becoming more elaborate. Understandably, an increasingly diverse background was being required from these professionals, mainly because data center technologies were repeatedly consolidating different areas of knowledge such as networking, storage, application, servers, cabling, and several others. And much to my chagrin, I had to admit that the job of creating an effective introduction to these technologies was getting even harder to “crack.”

After developing many learning road maps and customized trainings, I decided to challenge myself in writing a book that would address cutting-edge data center technologies and the core concepts they were based upon. From the start, the mammoth level of minutiae made me realize how close I was to a task such as writing a Beatles biography. And that exact thought inspired me to follow the steps of the best publications about the band: I had to use a *unifying theme*, something that could provide a firm backbone to a progressive presentation of these technologies. It was fairly easy for me to conclude that virtualization was this theme.

Nowadays, virtualization is deeply rooted in data center installations through technologies such as virtual memory, virtual gateways, VLANs, VRFs, virtual contexts, VDCs, vPCs, VNTag, VPLS, OTV, virtual LUNs, VSANs, IVR, NPV, FCoE, virtual machines, service profiles, virtual networking, virtual network services, and many others. All these successful techniques share a common characteristic: They were created to provide *resource optimization*. And for that reason, their examination opens up the chance to address the following:

- Traditional data center deployments and their limitations
- The benefits of each virtualization technology and their behavior
- The changes these technologies provide in data center designs and architectures

As the book cover suggests, virtualization has also modified the human aspects within data center environments. Relieved from the “chains of reality,” technical teams have been able to simplify operational tasks and accelerate the adoption of new IT models such as cloud computing. With such a central theme, it was just a question of defining *how* to approach it.

Goals and Methods

This book provides a gradual introduction to innovative data center technologies through a systematic examination of the main infrastructure virtualization techniques. And as an intentional outcome, the book also introduces fundamental concepts and definitions that are required from any professional who is involved with modern data center designs.

Because it is primarily focused on the three main data center infrastructure areas (networking, storage, and server), the book is not based on a single product nor is it written as a configuration guide. Instead, it leverages the broad Cisco Data Center portfolio (and other solutions from the Cisco ecosystem of partners and alliances) to analyze the behavior of each virtualization technique and to offer an architectural perspective of the virtualized data center.

Besides providing a technical account of the evolution of these areas, the book will address each virtualization technology through a flow of topics that involves

- A virtualization classification system (explained in the first chapter), which quickly informs the reader about the main characteristics of the specific technology
- A technology primer that immerses the reader in the physical challenges this virtualization technology overcomes
- A detailed analysis of the technique, including its characteristics, possibilities, scalability, results, and consequences
- A real-world use case scenario that demonstrates the examined technology “in action.”

I sincerely believe that design and deployment must be complementary processes. Therefore, *Data Center Virtualization Fundamentals* contains actual configuration examples that were exclusively created to illustrate each virtualization technology and its applicability to data center designs. Nonetheless, I have also included unusual topologies to specifically reinforce concepts explored throughout the book.

Who Should Read This Book

This book was written with a wide audience in mind. Because it provides an in-depth examination of data center virtualization technologies (from conceptualization to implementation), the book will satisfy beginners and experienced IT professionals alike.

In essence, its target audience comprises the following:

- Individuals with basic networking and operating system knowledge who are interested in modern data center design, deployment, and infrastructure optimization techniques
- Candidates for the Cisco Data Center certifications, including CCNA Data Center, CCNP Data Center, and CCIE Data Center

- Professionals that are specialized in a single data center technology area but also intend to acquire a broader architectural knowledge to accelerate their career development

How This Book Is Organized

With the explosion of information brought by the Internet, education in the twenty-first century must always present alternatives to the random accumulation of unstructured data. Therefore, I have intentionally applied constructivist learning theory principles (such as systematic analysis and concept synthesis) to distribute the content throughout the book. Although each chapter can be read out of sequence, their arrangement was designed to provide a logical progression of explanations that brings a more rewarding learning experience for the reader.

In times where blog posts and tweets provide “snacks” of information (do not get me wrong; there are nutritious knowledge bites out there), this book intends to serve a complete “meal,” where order and harmonization between chapters matter.

Chapters 1 through 17 and the appendixes cover the following topics:

- **Chapter 1, “Virtualization History and Definitions”:** This introductory chapter presents a historical account of virtualization in data center environments and, through some illustrative examples, provides a unified definition of virtualization in this context. It also proposes a classification system (which is called “virtualization taxonomy”) that will be used throughout the book to quickly introduce a new virtualization technology for the reader.
- **Chapter 2, “Data Center Network Evolution”:** Using the evolution of the Ethernet protocol as a canvas, this chapter addresses the main aspects and factors that govern traditional data center network topologies. It also discusses the general benefits that virtualization can offer to these networks.
- **Chapter 3, “The Humble Beginnings of Network Virtualization”:** Focused on the explanation of virtual local-area networks (VLAN) and Virtual Routing and Forwarding (VRF), this chapter provides a deep analysis of these well-established structures as virtualization techniques, illustrating the book approach and revealing important concepts that are hidden behind common knowledge.
- **Chapter 4, “An Army of One: ACE Virtual Contexts”:** This chapter discusses the importance of network services in data centers, concentrating on server load balancers. It presents virtual contexts as important tools that can increase flexibility and optimize hardware resources as these application environments scale.
- **Chapter 5, “Instant Switches: Virtual Device Contexts”:** The innovative characteristics of virtual device contexts (VDC) are detailed in this chapter, which also shows their applicability in challenging data center network scenarios.

- **Chapter 6, “Fooling Spanning Tree”:** This chapter demonstrates how virtualization techniques such as EtherChannel and virtual PortChannel (vPC) have adapted the limitation of Spanning Tree Protocol (STP) to the strict requirements of data center networks. It also introduces FabricPath, a Layer 2 multipathing technology that has provided the most secure path toward the replacement of STP in these environments.
- **Chapter 7, “Virtualized Chassis with Fabric Extenders”:** Fabric Extenders (FEX) constitute a virtualization technique that provides cabling optimization and network management consolidation in the data center network access layer. This chapter fully explores the many flavors of this technology.
- **Chapter 8, “A Tale of Two Data Centers”:** The classic problem of extending Layer 2 domains between geographically distinct data center sites is discussed throughout this chapter. It builds on concepts developed in previous chapters to offer a hands-on examination of the many different virtualization technologies that can solve this challenge.
- **Chapter 9, “Storage Evolution”:** This chapter explores the main concepts related to storage and storage access technologies that are used in data centers today. It also provides an account of how virtualization is deeply ingrained in the interpretation of stored data.
- **Chapter 10, “Islands in the SAN”:** Virtual storage-area networks (VSAN) can overcome Fibre Channel fabric challenges in a simple and elegant way. This chapter presents the necessary protocol concepts to understand how they can be applied in real-world scenarios.
- **Chapter 11, “Secret Identities”:** This chapter presents three virtualization techniques whose dissimulation tactics benefits data protection, environment isolation, and scalability in storage-area networks.
- **Chapter 12, “One Cable to Unite Us All”:** Binding concepts from network and storage virtualization, this chapter fully examines the details and benefits from the I/O consolidation brought about by Data Center Bridging (DCB) and Fibre Channel over Ethernet (FCoE).
- **Chapter 13, “Server Evolution”:** This chapter introduces the main concepts related to modern server architectures. It also presents server virtualization and describes how it has changed the operational landscape of data centers in the beginning of the twenty-first century. The chapter also deals with the definition of unified computing and explains how its innovative architecture principles can drastically simplify server environments.
- **Chapter 14, “Changing Personalities”:** Although server virtualization has helped to streamline server workloads within a data center, “bare metal” server provisioning and management are still considered massive challenges in these environments. This chapter demonstrates how service profiles can bring several server virtualization benefits to these scenarios.

- **Chapter 15, “Transcending the Rack”:** Demonstrating how the technologies explored in this book are extremely intertwined, this chapter shows how server virtualization has also revolutionized networking. It presents the virtual networking concepts through the analysis of VMware vSwitches, Nexus 1000V, and Virtual Machine Fabric Extender (VM-FEX).
- **Chapter 16, “Moving Targets”:** The way that virtual machines can migrate between different hosts and locations has also changed the way network services are deployed. This chapter explores the unique characteristics of services provided by solutions such as virtual firewalls (Virtual Security Gateway [VSG] and ASA 1000V), virtual accelerators (virtual Wide Area Application Services [vWAAS]), and virtual routers (CSR 1000V). It also presents site selection as a special network service and illustrates some solutions that can optimize client session routing to roaming virtual machines.
- **Chapter 17, “The Virtual Data Center and Cloud Computing”:** This chapter consolidates concepts explained throughout the book to discuss how 1+1 can be more than 2. It discusses how the deployment of multiple virtualization technologies has created a perfect storm for “cloud computing” momentum and how this IT delivery model is influencing the evolution of data center networks.
- **Appendix A, “Cisco Data Center Portfolio”:** To preserve the book’s focus on virtualization concepts and feature behavior, this appendix contains the description of all Cisco Data Center products that actually deploy these technologies.
- **Appendix B, “IOS, NX-OS, and Application Control Software Command-Line Interface Basics”:** If you are not used to the command-line interface characteristics from the different network operating systems used in this book, this appendix will introduce you to their most typical characteristics and definitions.

This page intentionally left blank

Virtualized Chassis with Fabric Extenders

“Form follows function – that has been misunderstood. Form and function should be one, joined in a spiritual union.” (Frank Lloyd Wright)

This chapter examines Fabric Extenders and how they optimize cabling and consolidate network management in data centers. It covers the following topics:

- Server Access Models
- Understanding Fabric Extenders
- Fabric Extender Topologies
- Use Case: Mixed Access Data Center

Table 7-1 categorizes Fabric Extenders using the virtualization taxonomy described in Chapter 1, “Virtualization History and Definitions.”

Table 7-1 *Fabric Extender Virtualization Classification*

| Virtualization Characteristics | Fabric Extender |
|---------------------------------------|--------------------------------------|
| Emulation | Modular Ethernet access switch |
| Type | Pooling |
| Subtype | Heterogeneous |
| Scalability | Hardware dependent ¹ |
| Technology Area | Networking |
| Subarea | Data, control, and management planes |

| Virtualization Characteristics | Fabric Extender |
|---------------------------------------|--|
| Advantages | Cabling savings, consolidated management |

¹ Refer to Appendix A, “Cisco Data Center Portfolio” for more details and the Cisco online documentation for updated information.

In data center projects, cabling is usually one of the first discussed topics. Unfortunately, it is very common that cabling design is defined before any network and server decisions are made. Because of its intrinsic affinity with three traditional data center teams (facilities, network, and server), server access architecture should be a shared decision that kick-starts optimized projects within each technology division.

To exacerbate this situation, one choice in particular does not seem to have a unanimous answer: Where do you position the access switches? Both Top-of-Rack and End-of-Row connectivity models present advantages and shortcomings, and during the last decades, they have attracted as many fans as detractors within the data center community.

Innovating again with its data center portfolio, Cisco introduced the concept of Fabric Extenders with the Nexus 2000 series. These devices are not Ethernet switches, but remote linecards of a virtualized modular chassis. This virtual entity permits server access to achieve the best aspects of both connectivity models.

Fabric Extenders allow scalable topologies that were not previously possible with traditional Ethernet switches. This chapter describes the design and configuration options of these devices and shows how they can optimize the server access layer within data centers.

Server Access Models

In 2005, the Electronic Industries Alliance (EIA) and the Telecommunications Industry Association (TIA) published the first formal specification for data center infrastructure: ANSI/TIA-942. This standard was intended to provide requirements and guidelines for the design and installation of a data center and includes the facility, network, and cabling design.

This specification defines horizontal cabling as the extension from the mechanical termination of the equipment distribution area (servers) to the horizontal distribution area (switches). Compared with backbone cabling (between switches and other communication equipment), horizontal cabling presents a much higher number of connections and, therefore, has a bigger impact over the entire infrastructure.

As discussed in Chapter 2, “Data Center Network Evolution,” the ANSI/TIA-942 specification supports the most popular server connectivity models: Top-of-Rack (ToR) and End-of-Row (EoR). These models define where the access layer switches are positioned in relation to the server localization, and consequently, how the horizontal cabling is designed.

In ToR-based networks, access switches (with 1 or 2 rack units) are usually installed at the top position inside the server cabinets. As a result, horizontal cabling is intrarack and can use a great variety of media such as twisted-pair, fiber, or twinax. To allow easier uplink upgrades, fiber is generally used in the redundant connections to upper-layer devices (core, aggregation, or edge routers).

Figure 7-1 depicts a Top-of-Rack access network.




Figure 7-1 *Top-of-Rack Network*

The ToR model permits

- Savings in horizontal cabling (because cable length is reduced)
- Provisioning of fully populated server cabinets
- Per-cabinet migration of connection technologies (Gigabit Ethernet to 10 Gigabit Ethernet, for example).

The number of servers per cabinet heavily influences the port utilization in Top-of-Rack designs. ToR switches usually have 24, 32, 48, or 96 Ethernet interfaces, and if a cabinet does not support a considerable number of servers, some switch interfaces can remain unused.

To increase port utilization on low-populated cabinets, ToR designs can be adapted to other horizontal cabling variations. Figure 7-2 portrays an example where each switch is the “**Top-of-Many-Racks.**”

Note Data center power distribution and cooling capacity usually define how many servers can be installed per cabinet.

A pointed disadvantage for the ToR model is the management effort that must be spent on multiple switches. Because a data center can span hundreds or even thousands of server racks, regular operations (such as the identification of interfaces and firmware upgrades) can become quite challenging.

Alternatively, the EoR model allows the connectivity management of hundreds of servers, which are installed in a row of cabinets, with a pair of devices.




Figure 7-2 “Top-of-Many-Racks” Network

EoR switches are typically modular chassis with UTP-based interfaces for horizontal cabling and multiple fiber connections to the network upper-layer devices (core, aggregation, or edge routers).

Figure 7-3 illustrates an End-of-Row access network, while Figure 7-4 portrays a variation called **Middle-of-Row**. The latter decreases the average cable length, which can achieve substantial cost savings depending on the number of cabinets in each row.




Figure 7-3 End-of-Row




Figure 7-4 Middle-of-Row

EoR topologies are more flexible for low-density cabinets, because the horizontal cabling can potentially reach any rack in the row, leveraging the higher number of ports of modular chassis switches.

Cabling sprawl is the main disadvantage of EoR topologies. In a data center, excessive EoR cabling brings consequences such as

- Difficult cable management, troubleshooting, and decommission
- Blockage for air cooling (if it is installed underneath a raised access floor)

Data centers usually invest in structured cabling, preprovisioning a large part of the connections to avoid excessive cable installation. In these situations, EoR-based networks can only adopt new connectivity technologies if they were previously considered in the original cabling design.

Over the last decades, none of the access models were clearly defined as the “best” for cabling implementations. A design decision between ToR or EoR models depends on technical experience, knowledge about the server environment, and investment resources. However, best results should not be expected when there is little interaction between the facilities, network, and server teams.

Understanding Fabric Extenders

In 2009, Cisco launched the Nexus 2000 Fabric Extender series. Not conceived to be used as standalone Ethernet switches, these devices are remote linecards that are managed by a *parent* switch, such as a Nexus 5000, Nexus 6000, UCS Fabric Interconnect, or Nexus 7000 (with appropriate modules).

Fabric Extenders (or FEX) enable a data center to leverage the advantages from both ToR and EoR models because

- Multiple Fabric Extenders can be managed from a single parent switch (similar to EoR).
- A Fabric Extender can be installed inside a server cabinet and decrease cabling costs (similar to ToR).

Figure 7-5 illustrates an example of a server access topology that uses Fabric Extenders.




Figure 7-5 *Fabric Extender Topology*

As the figure suggests, a parent switch and multiple Fabric Extenders are elements of a virtualized modular chassis. Inside this virtual structure, every management operation is performed on the parent switch (which performs the role of a supervisor module of such

chassis), and Ethernet frames are exchanged on the Fabric Extender interfaces (which represent the chassis interface modules).

The main endeavor of Fabric Extenders is to keep the configuration complexity within the parent switches and drive simplicity toward the server interfaces.

The Fabric Extender architecture introduces new types of ports to the network, including physical and virtual interfaces. Figure 7-6 deals with the interaction between the following FEX-related interfaces:

- **Fabric Interface (FIF):** This is a physical interface created to connect the Fabric Extender to a parent switch. It cannot be used for any other purpose.
- **Host Interface (HIF):** This is a standard physical Ethernet interface designed for server connection.
- **Logical Interface (LIF):** This is a data structure in the parent switch that emulates an Ethernet interface. It carries properties such as VLAN membership, access control list (ACL) labels, and STP states and is mapped to a virtual interface created on a Fabric Extender.
- **Virtual Interface (VIF):** This is a logical entity inside Fabric Extenders that receives its configuration from the parent switch, and it is used to map frames to a switch Logical Interface (LIF).




Figure 7-6 *Fabric Extender Interfaces*

The mapping between a parent switch LIF and a Fabric Extender VIF is called a *Virtual Network Link (VN-Link)*, and it is defined through a special tag that is inserted on all Ethernet frames that traverse these physical links. This extra header is called a *Virtual*

Network Tag (VNTag), and its main objective is to differentiate frames received from (or sent to) distinct FEX host interfaces.

Figure 7-7 depicts the 6-byte VNTag header implemented on Cisco Fabric Extenders.

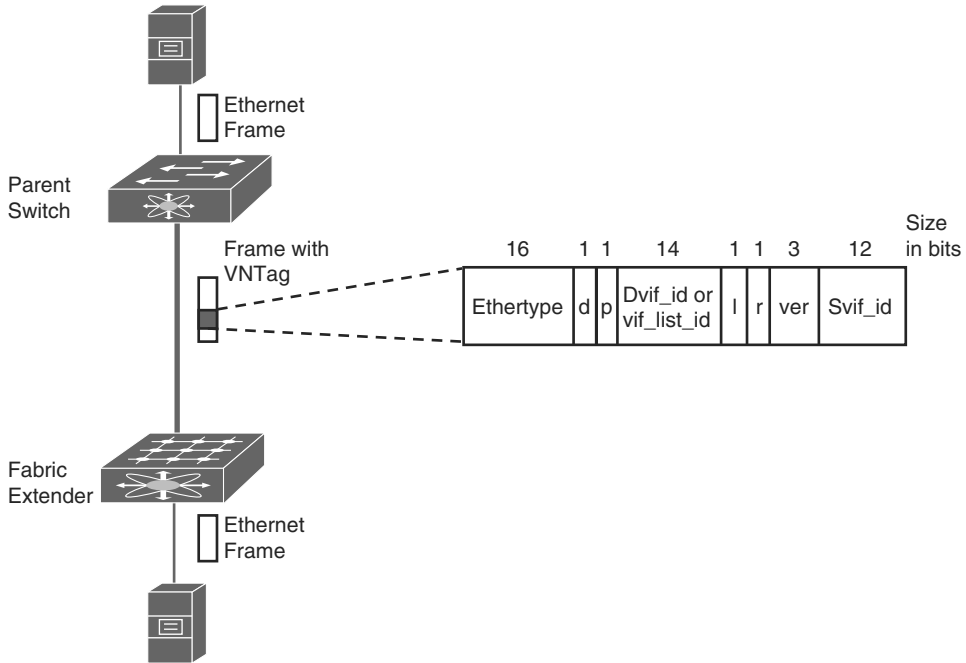


Figure 7-7 Virtual Network Tag

The VNTag is inserted between the source MAC address and the IEEE 802.1Q fields from the original Ethernet frame. The VNTag fields are

- **Ethertype:** This field identifies a VNTag frame. IEEE reserved the value 0x8926 for Cisco VNTag.
- **Direction bit (d):** A 0 indicates that the frame is traveling from the FEX to the parent switch. A 1 means that the frame is traveling from the parent switch to the FEX.
- **Pointer bit (p):** A 1 indicates that a Vif_list_id is included in the tag. A 0 signals that a Dvif_id is included in the frame.
- **Virtual Interface List Identifier (Vif_list_id):** This is a 14-bit value mapped to a list of host interfaces to which this frame must be forwarded.
- **Destination Virtual Interface Identifier (Dvif_id):** This is a 12-bit value mapped to a single host interface to which an Ethernet frame will be forwarded.

- **Looped bit (l):** This field indicates a multicast frame that was forwarded out the switch port and later received. In this case, the FEX checks the Svif_id and filters the frame from the corresponding port.
- **Reserved bit (r):** This bit is reserved for future use.
- **Version (ver):** This value is currently set to 0. It represents the version of the tag.
- **Source Virtual Interface Identifier (Svif_id):** This is a 12-bit value mapped to the host interface that received this frame (if it is going from the FEX to the parent switch).

When an Ethernet frame is received on a host interface,

The Fabric Extender adds a VNtag to the frame and forwards it to one of the fabric interfaces.

The parent switch recognizes the logical interface that sent the frame (through the Svif_id field), removes the tag, and forwards it according to its MAC address table.

Tip The parent switch MAC address table is updated if an unknown MAC address originates the frame received on an HIF. The new entry points to the logical interface index.

In the other direction (parent switch receives a frame that is destined to a FEX host interface),

The parent switch reads the frame destination MAC address and forwards it to a logical interface index in its MAC address table.

The switch inserts the VNtag associated with the logical interface and forwards it to the correct FEX.

Receiving this frame, the FEX recognizes the associated VIF (through the Dvif_id), removes the VNtag, and sends it to the mapped host interface.

From a data plane perspective, forwarding in Fabric Extender host interfaces completely depends on the parent switch. Consequently, a frame exchange between two host interfaces always traverses the parent switch, even if they are located in the same Fabric Extender.

Note Because of its hardware simplicity, Nexus 2000 Fabric Extenders have a very low latency (as low as 500 nanoseconds at the time of this writing) when compared to traditional Ethernet switches.

Multidestination VNtag frames (such as flooding, broadcast, and multicast) are characterized with the Pointer bit and Vif_list_id fields. When such a frame is being forwarded

from the parent switch to the Fabric Extender, it is replicated inside the FEX to reduce the fabric interface traffic.

Note Cisco has been shipping VNTag as a prestandard port extension protocol since 2009. This additional header is intended to deliver the same capabilities defined in the IEEE 802.1BR (formerly IEEE 802.1Qbh) standard, which was published in July 2012. At the time of this writing, Cisco is expected to deliver products that fully support a VNTag and standards-based solutions.

Fabric Extender Options

Cisco has launched several Fabric Extender models since the Nexus 2148T was first shipped in 2009. Subsequent Cisco Fabric Extender products differ from each by factors such as

- **Fabric interfaces:** Most Cisco Fabric Extenders have four or eight 10 Gigabit Ethernet interfaces, which can use fiber (10GBASE-ER, 10GBASE-LR, Fabric Extender transceivers) or twinax cables. At the time of this writing, the Nexus 2248PQ exclusively supports 4 QSFP+ fabric interfaces that can deploy four 40-gigabit Ethernet connections or, using breakout cables, sixteen 10-gigabit Ethernet connections.
- **Host interfaces:** 8, 24, 32, or 48 Ethernet ports, which can be fixed (1000BASE-T, 100BASE-TX/1000BASE-T, 10GBASE-T) or SFP based (10GBASE-ER, 10GBASE-LR, twinax, 1000BASE-T).
- **Memory buffer:** Some models have larger shared buffers for applications such as large databases, shared storage, and video editing.
- **Form factor:** Fabric Extenders can be accommodated in server cabinets or in select blade server chassis (the latter option includes the UCS I/O Module).
- **Network capabilities:** These factors include multiple PortChannel members, number of quality of service (QoS) queues, and Fibre Channel over Ethernet (FCoE).
- **Choice of airflow:** Some models offer the choice to alternatively deploy front-to-back or back-to-front airflow.

A virtualized chassis can have different Fabric Extenders depending on server access requirements. This flexibility facilitates technology evolution, because migrations can be executed per server cabinet or even per server connection.

Devices as the Nexus 7000, Nexus 6000, Nexus 5000, and UCS Fabric Interconnect can act as parent switches to Fabric Extenders. Each one of these supports distinct Fabric Extender features and has different scalability characteristics.

Note For more details about Cisco Fabric Extenders models (and their parent switches), you can read Appendix A, “Cisco Data Center Portfolio.” If you want to verify the capabilities supported on your hardware and software combination, also refer to the Cisco online documentation for the most recent information.

Connecting a Fabric Extender to a Parent Switch

When there is one active connection between them, a Fabric Extender and its parent switch use Satellite Discovery Protocol (SDP) periodic messages to discover each other. After this formal introduction, the Fabric Extender deploys a Satellite Registration Protocol (SRP) request to register itself to the parent switch.

Figure 7-8 illustrates the FEX discovery and registration process when a switch interface is configured with the `switchport mode fex-fabric` command.

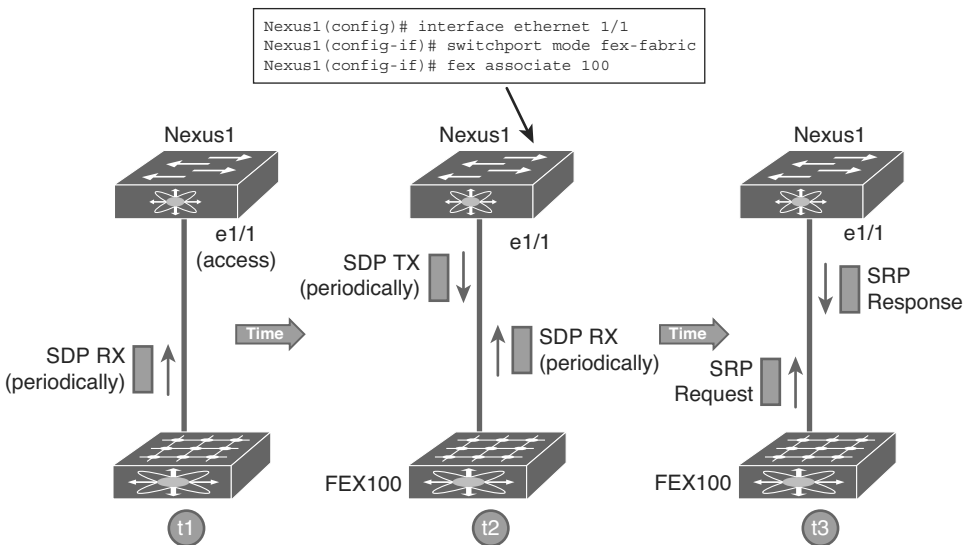


Figure 7-8 Fabric Extender Discovery and Registration Process

Example 7-1 shows a sample of the SDP messages sent from the Fabric Extender before any FEX configuration is executed in Nexus1 (at instant t1). As expected, the Fabric Extender automatically sends SDP messages as soon as it has an active fabric interface.

Example 7-1 SDP Frames Sent at t1

```
! Enabling the Fabric Extender processes in Nexus1
Nexus1# feature fex
! Enabling the tracing of SDP and SRP packets
Nexus1# debug fex pkt-trace
```

```

! Sample SDP Message received at interface Ethernet 1/1 (and repeated every 3
seconds)
08:35:03.086490 fex: Sdp-Rx: Interface: Eth1/1, Fex Id: 0, Ctrl Vntag: -1, Ctrl
Vlan: 1
08:35:03.086511 fex: Sdp-Rx: Refresh Intvl: 3000ms, Uid: 0x80a3e61d9cc8, device:
Fex, Remote link: 0x20000080
08:35:03.086523 fex: Sdp-Rx: Vendor: Cisco Systems Model: N2K-C2232PP-10GE Serial:
JAF1509ECGT
[output suppressed]

```

In these SDP messages, the Fabric Extender exposes the VLAN from which it expects to receive control commands, SDP refresh interval (3 seconds), and hardware information.

After interface e1/1 in Nexus1 is configured with the **switchmode fex-fabric** and **fex associate** commands (instant t2 in Figure 7-8), this interface starts to send SDP packets, and both devices discover each other.

Example 7-2 details the Fabric Extender discovery process from the parent switch perspective.

Tip Fabric Extender identifiers must belong to the 100–199 range.

Example 7-2 SDP Frames Sent After Ethernet 1/1 Configuration

```

! SDP Messages sent from interface Ethernet 1/1 with Control VNTag -1 and Control
VLAN 4042
08:36:27.007428 fex: Sdp-Tx: Interface: Eth1/1, Fex Id: 100, Ctrl Vntag: -1, Ctrl
Vlan: 4042
08:36:27.007445 fex: Sdp-Tx: Refresh Intvl: 3000ms, Uid: 0xc03cc7730500, device:
Switch, Remote link: 0x1a000000
2012 Jun 4 08:36:27.007456 fex: Sdp-Tx: Vendor: Model: Serial: -----
[output suppressed]
! SDP Messages sent from interface Ethernet 1/1 with new VNTag (zero)
08:36:27.101688 fex: Sdp-Tx: Interface: Eth1/1, Fex Id: 100, Ctrl Vntag: 0, Ctrl
Vlan: 4042
08:36:27.101704 fex: Sdp-Tx: Refresh Intvl: 3000ms, Uid: 0xc03cc7730500, device:
Switch, Remote link: 0x1a000000
08:36:27.101715 fex: Sdp-Tx: Vendor: Model: Serial: -----
! Notice how the SDP Messages received from the FEX have changed
08:36:27.125135 fex: Sdp-Rx: Interface: Eth1/1, Fex Id: 0, Ctrl Vntag: 0, Ctrl Vlan:
4042
08:36:27.125153 fex: Sdp-Rx: Refresh Intvl: 3000ms, Uid: 0x80a3e61d9cc8, device:
Fex, Remote link: 0x20000080
08:36:27.125164 fex: Sdp-Rx: Vendor: Cisco Systems Model: N2K-C2232PP-10GE Serial:
JAF1509ECGT
[output suppressed]

```

After the discovery is complete, the Fabric Extender sends an SRP Request message and waits for an SRP Response from the parent switch (instant t3). The registration process completes the FEX detection on the parent switch, and SDP messages continue to be exchanged between both devices after the registration.

Example 7-3 shows the registration process and describes how the Fabric Extender becomes operational for Nexus1.

Example 7-3 SRP Registration and FEX Becoming Operational

```

! FEX sens an SRP Request
08:36:30.114980 fex: Srp Req: Interface: Eth1/1, Uid: 0x80a3e61d9cc8, Card Id: 82,
IPC ver: 21
08:36:30.114992 fex: Srp Req: Version: 5.1(3)N2(1a), Interim Version: 5.1(3)N2(1a)
! Nexus1 sends an SRP Response confirming that FEX firmware is
supported . . .
08:36:30.116726 fex: Srp Resp: Interface: Eth1/1, Fex id: 100 Ver Chk: Compatible,
Img Uri:
! . . . and assigns an internal IP address for the FEX
08:36:30.116744 fex: Srp Resp: MTS addr: 0x2102, IP addr: 127.15.1.100/0 Switch MTS:
0x101, Switch Ip: 127.15.1.250
! After that the parent switch has a new remote linecard
08:26:16 5548P-3-174 %$ VDC-1 %$ %SATCTRL-FEX100-2-SOHMS_ENV_ERROR: FEX-100 Module
1: Check environment alarms.
[output suppressed]
08:36:33 5548P-3-174 %$ VDC-1 %$ %PFMA-2-FEX_STATUS: Fex 100 is online

```

From this moment on, the Fabric Extender is a remote linecard for the parent switch's supervisor module. Example 7-4 shows how the FEX interfaces can be seen and controlled on Nexus1.

Example 7-4 Verifying the Fabric Extender Status

```

! Verifying the operational Fabric Extenders
Nexus1# show module fex
FEX Mod Ports Card Type                               Model                Status.
-----
100 1    32    Fabric Extender 32x10GE + 8x10G N2K-C2232P-10GE present
[output suppressed]
! Verifying the available FEX interfaces
Nexus1# show interface brief
[output suppressed]
-----
Ethernet    VLAN    Type Mode    Status Reason                Speed Port
Interface                                     Ch #
-----

```

```
! Notice that all interfaces have the FEX ID and that interface e100/1/1 is down
Eth100/1/1    1  eth  access down    SFP not inserted 10G(D) --
Eth100/1/2    1  eth  access down      SFP not inserted 10G(D) --
[output suppressed]
Eth100/1/32   1  eth  access down      SFP not inserted 10G(D) --
```

Nexus1 logical interfaces (LIF) can be recognized through the FEX ID/slot/port index format. And because these interfaces are within Nexus1 configuration reach, its default configuration is applied to the FEX interfaces as well (**switchport mode access** and **no shutdown** commands, in this case).

With a control VLAN (4042) and an internal IP address (127.15.1.100), the Fabric Extender is configured through the Virtual Interface Configuration (VIC) protocol after its operationalization. Based on command/response messages, this protocol is also responsible for configuring the forwarding tables inside each FEX. The parent switch uses the VIC protocol to assign the Dvif_id to each virtual interface in the FEX and retrieve the FEX virtual interface list identifiers (Vif_list_id) for multidestination traffic.

Note Chapter 15, “Transcending the Rack,” discusses a VNTag-based technology called Virtual Machine Fabric Extender (VM-FEX), where specialized adapters use the VIC protocol to dynamically request the creation, deletion, enabling, and disabling of logical interfaces in the parent switch.

Fabric Extended Interfaces and Spanning Tree Protocol

In the previous section, you learned that FEX host interfaces can be visualized as any other interface on the parent switch. These interfaces can also inherit various features from the same switch, justifying the concept of the virtualized modular chassis.

Example 7-5 illustrates the enablement of interface Ethernet 100/1/1, which is using a 100BASE-T Small Form Pluggable (SFP) transceiver.

Example 7-5 FEX Interface Configuration

```
! Configuring interface Ethernet 100/1/1 to accept the GigabitEthernet SFP
Nexus1(config)# interface ethernet 100/1/1
Nexus1(config-if)# speed 1000
! Verifying the interface status
Nexus1(config-if)# show interface ethernet 100/1/1 brief
```

| Ethernet Interface | VLAN | Type | Mode | Status | Reason | Speed | Port Ch # |
|--------------------|------|------|------|--------|--------|-------|-----------|
| Ethernet Interface | | | | | | | |

```
-----
Eth100/1/1    1          eth  access up    none          1000 (D)  --
```

Nevertheless, there are some differences on how a host interface behaves in the Spanning Tree Protocol context. In Example 7-6, you can spot three commands that expose this distinct behavior.

Example 7-6 Fabric Extender Interface and Spanning Tree Protocol

```
! Verifying the STP status on the Fabric Interface
Nexus1# show spanning-tree interface ethernet 1/1
No spanning tree information available for Ethernet1/1
! Verifying the STP status on FEX interface
Nexus1# show spanning-tree interface ethernet 100/1/1

Vlan                Role Sts Cost      Prio.Nbr Type
-----
VLAN0001            Desg FWD 4        128.1025 Edge P2p
! Verifying the BPDU features on the FEX interface
Nexus1# show spanning-tree interface e100/1/1 detail | include Bpdu
    Bpdu guard is enabled
    Bpdu filter is enabled by default
Nexus1# show spanning-tree interface e100/1/1 detail | include BPDU
    BPDU: sent 11, received 0
```

From the example, it is possible to infer that

The interface connected to the Fabric Extender (Ethernet 1/1) does not participate in STP, because from the Nexus1 perspective, it is considered a “backplane connection.”

The host interfaces on the FEX are configured as RSTP edge point-to-point ports.

By default, host interfaces are configured with bridge protocol data unit (BPDU) filter and BPDU guard. These configurations mean, respectively, that the interfaces will not process any received BPDUs, and will actually be disabled if such a frame arrives at them. Nevertheless, as a safeguard mechanism, these interfaces do send some BPDUs when they are activated to avoid loops that can be caused with a mistakenly direct connection between two host interfaces.

As a conclusion, the parent switch controlling Fabric Extenders as linecards forms a virtualized chassis designed for host connections. In my opinion, a correct characterization for the virtualization technique detailed in this chapter would be *virtualized modular access switch*.

Caution Although Fabric Extenders are primarily designed for server connections, they also support switches connected to host interfaces as long as you disable STP on these devices. In consequence, to eliminate loops in redundant active connections, the switches must deploy a redundancy mechanism that is not STP dependent (Cisco Flex Link, for example).

Fabric Interfaces Redundancy

A Fabric Extender with multiple connections to a parent switch can have two different behaviors in the case of a fabric interface failure. By default, a FEX deploys an interface pinning policy to control which host interfaces are shut in the case of such a failure.

The objective of static pinning is to keep the exact amount of bandwidth at the host ports in the case of a fabric interface failure. Hence, for the remaining available host interfaces, there is no oversubscription increase.

Example 7-7 shows how a fabric interface (Ethernet 1/1) is “pinned” to an active host interface (Ethernet 100/1/1).

Example 7-7 Static Pinning

```
Nexus1# show fex 100 detail
FEX: 100 Description: FEX0100 state: Online
[output suppressed]
Fabric interface state:
  Eth1/1 - Interface Up. State: Active
  Eth1/2 - Interface Up. State: Active
! Interface Ethernet 100/1/1 is active as long as Ethernet 1/1 is active
Fex Port      State Fabric Port
  Eth100/1/1  Up    Eth1/1
  Eth100/1/2  Down  None
[output suppressed]
  Eth100/1/32 Down  None
```

The example output exhibits the host interface status dependence whether or not its associated fabric interface is active. If Ethernet 1/1 fails, Ethernet 100/1/1 is automatically disabled, and hopefully, the connected server should activate a fault tolerance mechanism (NIC teaming).

Example 7-8 portrays the host interface behavior when the fabric interface is disabled. Notice that Ethernet 100/1/1 fails too, ignoring the fact that another fabric interface (Ethernet 1/2) is still active.

Example 7-8 *Failure in Static Pinning*

```

! Disabling fabric interface e1/1
Nexus1(config)# interface ethernet 1/1
Nexus1(config-if)# shutdown
! Vefifying e100/1/1 status
Nexus1(config-if)# show fex 100 detail
FEX: 100 Description: FEX0100 state: Online
[output suppressed]
Pinning-mode: static Max-links: 1
[output suppressed]
Fabric interface state:
  Eth1/1 - Interface Down. State: Configured
  Eth1/2 - Interface Up. State: Active
Fex Port      State Fabric Port
  Eth100/1/1  Down   Eth1/1
  Eth100/1/2  Down   None
[output suppressed]
  Eth100/1/32 Down   None

```

The **pinning max-link** command divides the host interfaces among a maximum number of fabric interfaces. With our 32-port Fabric Extender, if you issue the **pinning max-link 4** command, eight interfaces will be disabled if a fabric port fails.

In static pinning, host interface assignment follows the fabric interface order of configuration. However, if you reload the parent switch or execute the **fex pinning redistribute** command, the host interface groups will be reassigned to the fabric interface numerical order.

Table 7-2 illustrates this exact scenario, where the order of configuration for the fabric interfaces was Ethernet 1/3, Ethernet 1/2, Ethernet 1/4, and Ethernet 1/1.

Table 7-2 *Fabric Extender Pinning Example*

| Parent Switch Interface | Associated Host Interfaces (Following Order of Configuration) | Associated Host Interfaces (After Reload or Redistribution) |
|-------------------------|---|---|
| e1/1 | 25 to 32 | 1 to 8 |
| e1/2 | 9 to 16 | 9 to 16 |
| e1/3 | 1 to 8 | 17 to 24 |
| e1/4 | 17 to 24 | 25 to 32 |

Tip Both the `pinning max-link` and `fex pinning redistribute` commands are disruptive to the Fabric Extender host interfaces.

Alternatively, the configuration of a PortChannel between the parent switch and the Fabric Extender guarantees that every host interface remains operational if a fabric interface fails. However, the remaining bandwidth to the parent switch will be shared by all host ports (increasing the oversubscription).

Example 7-9 presents the results from a PortChannel configuration on the fabric interfaces (shown later in Figure 7-9).

Example 7-9 *PortChannel Configuration*

```
! Verifying how FEX 100 uses PortChannel 100 as its fabric interface
Nexus1# show fex 100 detail
FEX: 100 Description: FEX0100 state: Online
[output suppressed]
! When using PortChannel Max-links must be one
Pinning-mode: static Max-links: 1
[output suppressed]
Fabric interface state:
  Po100 - Interface Up. State: Active
  Eth1/1 - Interface Up. State: Active
  Eth1/2 - Interface Up. State: Active
Fex Port      State Fabric Port
  Eth100/1/1  Up      Po100
  Eth100/1/2  Down    None
[output suppressed]
  Eth100/1/32 Down    None
```

When fabric interfaces are aggregated, the max-link pinning must be set to 1 (because it is not possible to have more than one upstream PortChannel per FEX), and the interface must be configured in mode ON because Fabric Extenders do not support Link Aggregation Control Protocol (LACP).

Example 7-10 demonstrates that a failure in Ethernet 1/1 does not have any effect over interface Ethernet 100/1/1 when the fabric interfaces are aggregated. As expected, PortChannel 100 remains operational if a single member interface is active (Ethernet 1/2).

Example 7-10 *Failure in Ethernet 1/1*

```
! Disabling fabric interface e1/1
Nexus1(config)# interface ethernet 1/1
Nexus1(config-if)# shutdown
```



```

! Vefifying e100/1/1 status
Nexus1(config-if)# show fex 100 detail
FEX: 100 Description: FEX0100 state: Online
[output suppressed]
Pinning-mode: static Max-links: 1
[output suppressed]
Fabric interface state:
  Po100 - Interface Up. State: Active
  Eth1/1 - Interface Down. State: Configured
  Eth1/2 - Interface Up. State: Active
! PortChannel is still operational
Fex Port      State Fabric Port
  Eth100/1/1  Up      Po100
  Eth100/1/2  Down    None
[output suppressed]
  Eth100/1/32 Down    None
    
```

Figure 7-9 exhibits the behavior for both fabric interface redundancy options (static pinning and PortChannel) and details the configuration required in each one.

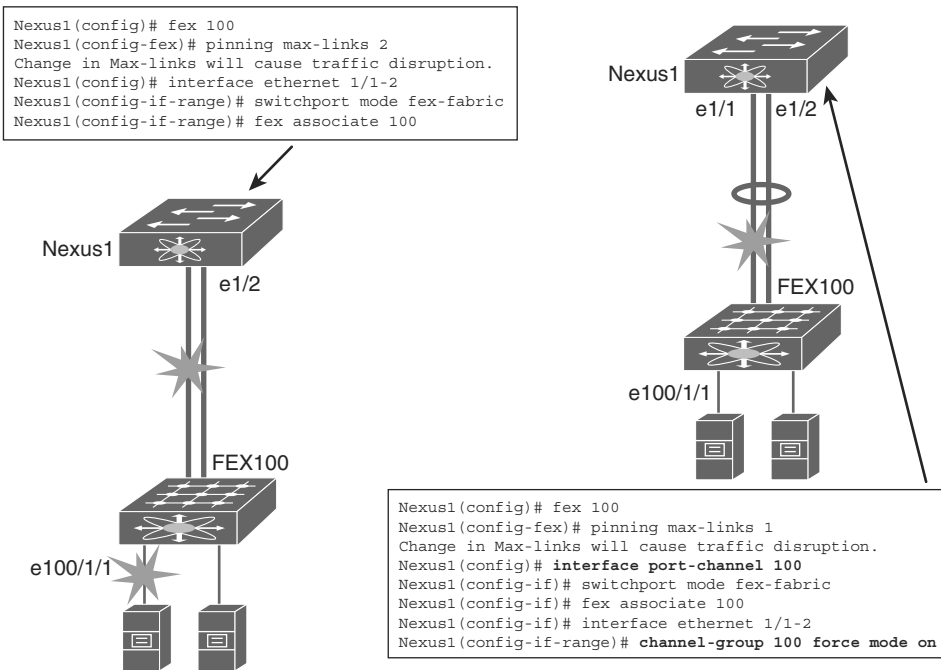


Figure 7-9 Static Pinning Versus PortChannel: Configuration and Host Interface Behavior

Fabric Extender Topologies

In the previous section, you were presented with simple scenarios that demonstrated basic functionalities of a Fabric Extender connected to one parent switch. Nevertheless, it is paramount that you understand the principles behind highly available FEX topologies.

There are basically two classes of topologies that provide fault tolerance in Fabric Extender designs:

Straight-through: Where a Fabric Extender is connected to a single parent switch

Dual-homed: Where a Fabric Extender is connected to a pair of parent switches

An example of each of these topologies is presented in Figure 7-10.

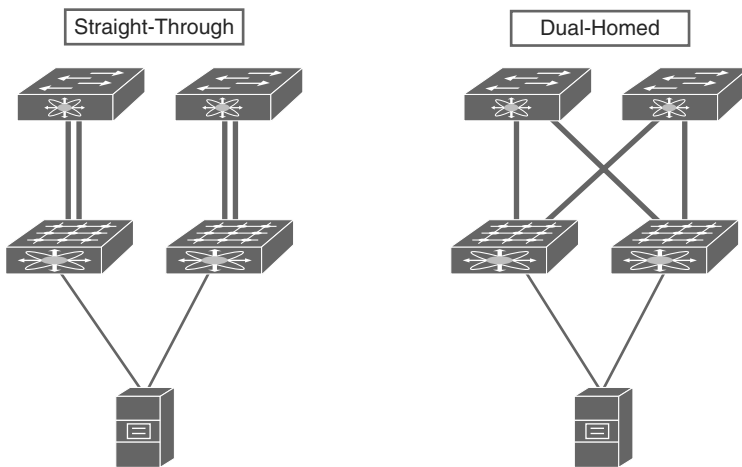


Figure 7-10 *Fabric Extender Topology Types*

Note In the next sections, I will discuss topologies whose support varies depending on the parent switch and Fabric Extender hardware models and software versions. Because of the NX-OS constant state of evolution, I will avoid transcribing a matrix of supported topologies that will be quickly outdated. I highly recommend that you refer to the Cisco online documentation for the most recently supported FEX topologies.

Straight-Through Topologies

In straight-through topologies, it is recommended that each host has interfaces connected to Fabric Extenders that are managed by distinct parent switches. This practice avoids the total of loss of connectivity for a server in case of a switch failure.

Straight-through designs create a pair of NX-OS virtualized chassis with a single supervisor module in each. Likewise, IEEE 802.3ad-compatible servers can leverage active-active connections using virtual PortChannels (vPCs).

Figure 7-11 demonstrates how vPCs can be deployed in a straight-through topology comprised of two switches (Nexus1 and Nexus2) and two Fabric Extenders (FEX110 and FEX120). In the figure, each Fabric Extender uses two aggregated fabric connections to its parent switch. Both server interfaces are connected to access interfaces in VLAN 50 (e110/1/1 in Nexus1 and e120/1/1 in Nexus2) which will be vPC 10 member ports.

Tip Notice that Nexus1 is configured to be the primary vPC peer, using the role **priority 1** command.

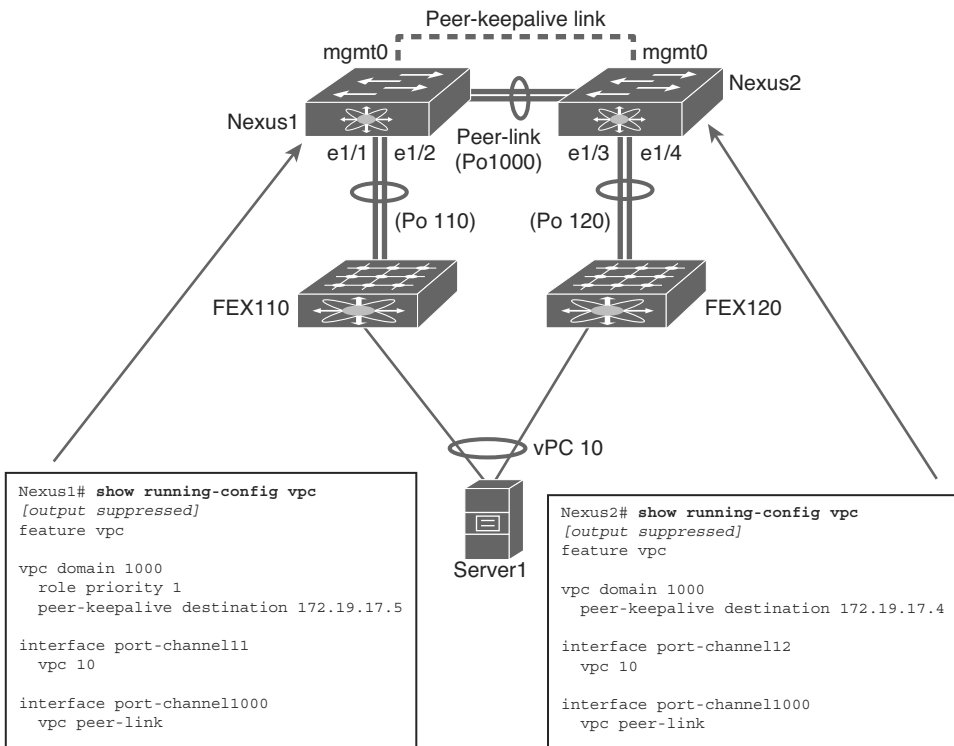


Figure 7-11 Virtual PortChannel Straight-Through Topology

For clarification, the PortChannel configuration options in this topology are explained in Table 7-3.

Table 7-3 *PortChannel Options*

| Switch | PortChannel ID | Interfaces | Remote Peer | LACP |
|--------|----------------|--------------|-------------|-----------------|
| Nexus1 | 11 (vPC 10) | e110/1/1 | Server A | Depends on peer |
| Nexus1 | 110 | e1/1, e1/2 | FEX110 | No |
| Nexus1 | 1000 | e1/31, e1/32 | Nexus2 | Depends on peer |
| Nexus2 | 12 (vPC 10) | e120/1/1 | Server A | Depends on peer |
| Nexus2 | 120 | e1/3, e1/4 | FEX120 | No |
| Nexus2 | 1000 | e1/31, e1/32 | Nexus1 | Depends on peer |

Note PortChannels 110 and 120 are not mandatory for this configuration because static pinning would work as well.

In Examples 7-11 and 7-12, you can verify the status of the virtual PortChannel created to aggregate Server1 Ethernet interfaces.

Example 7-11 *Virtual PortChannel 10 Status in Nexus1*

```
! Verifying vPC 10 status
Nexus1# show vpc 10
vPC status
-----
id      Port      Status Consistency Reason      Active vlans
-----
10      Po10     up      success      success      50
```

Example 7-12 *Virtual PortChannel Status in Nexus2*

```
! Verifying vPC 10 status
Nexus1# show vpc 10
vPC status
-----
id      Port      Status Consistency Reason      Active vlans
-----
10      Po10     up      success      success      50
```

In a straight-through vPC topology, the consequences of peer link total failure are the same as expected in a standard vPC configuration: All vPC member ports in the secondary vPC peer are automatically disabled. Example 7-13 exhibits this situation in Nexus2, when PortChannel 1000 is disabled on Nexus1 (through a **shutdown** command in the PortChannel).

Example 7-13 *Peer Link Effect on Nexus2*

```
! After the eer-link PortChannel is shut down in Nexus1, a peer-link failure
detected in Nexus2
Nexus2#
17:02:51 Nexus2 %$ VDC-1 %$ %VPC-2-VPC_SUSP_ALL_VPC: Peer-link going down, suspending
all vPCs on secondary
! Inspecting which interfaces are down due to peer-link failure
Nexus2# show interface brief | include vpc
Eth120/1/1    50    eth access down vpc peerlink is down 1000(D) 10
! Verifying FEX status in Nexus2
Nexus2# show fex
  FEX          FEX          FEX          FEX
Number  Description      State          Model          Serial
-----
120     FEX0120          Online        N2K-C2232PP-10GE  SSI150606H2
```

As expected, even though FEX120 is still operational, only vPC member ports are disabled because of the peer link failure.

Orphan ports, connected only to the secondary peer, can also be isolated in straight-through vPC topologies. The same recommendations I have described in Chapter 6, “Fooling Spanning Tree,” apply for these interfaces.

Figure 7-12 depicts other possible straight-through topologies available at the time of this writing.

Tip I usually recommend straight-through topologies for scenarios that require higher access port scalability (because they can deploy a twofold increase in the number of Fabric Extenders that a single parent switch supports).

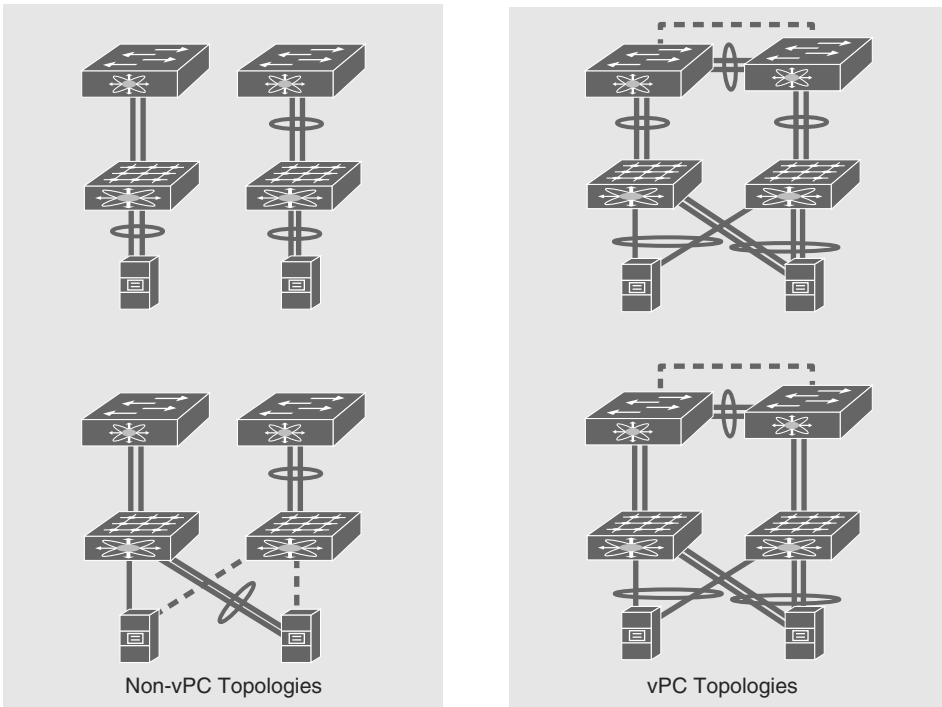


Figure 7-12 *Straight-Through Topologies*

Dual-Homed Topologies

In a classical modular Ethernet switch, redundant supervisor modules can control the installed linecards. In a virtualized access switch, the same redundancy can be achieved with the connection of two parent switches to a single Fabric Extender. However, you should not expect that the straightforward FEX configuration is sufficient to build a dual-homed topology: By default, a Fabric Extender can only be managed by a single parent switch.

To illustrate this behavior, consider the topology and configuration detailed in Figure 7-13. There, the Fabric Extender is connected and registered to Nexus1, and afterward, connected to Nexus2.

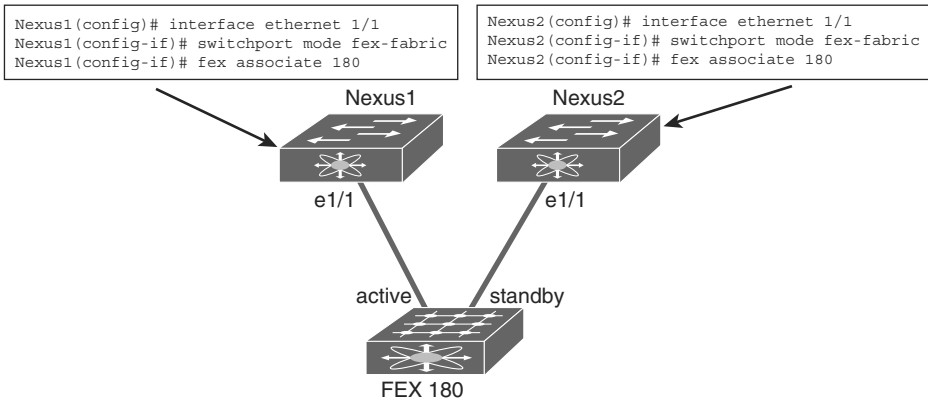


Figure 7-13 Active-Standby Dual-Homed Topology

Examples 7-14 and 7-15 demonstrate how each switch detects FEX180.

Example 7-14 Fabric Extender Connected on Two Parent Switches

```
! Verifying FEX status in Nexus1
Nexus1# show fex
```

| FEX Number | FEX Description | FEX State | FEX Model | FEX Serial |
|------------|-----------------|-----------|------------------|-------------|
| 180 | FEX0180 | Online | N2K-C2232PP-10GE | SSI150606J3 |

Example 7-15 Single FEX Connected on Two Parent Switches

```
! Verifying FEX status in Nexus2
Nexus2# show fex
```

| FEX Number | FEX Description | FEX State | FEX Model | FEX Serial |
|------------|-----------------|-----------|------------------|-------------|
| --- | ----- | Connected | N2K-C2232PP-10GE | SSI150606J3 |

The Fabric Extender remains in the **Connected** state (and not **Online**) in Nexus2 because it is already registered to Nexus1. Therefore, Nexus2 and the Fabric Extender only discover each other, without advancing to the registration phase.

If the connection to Nexus1 fails, the Fabric Extender registers itself to Nexus2 and remains in that state (even if the connection to Nexus1 is reactivated).

However, after the parent switch transition, the host interfaces would remain unconfigured. The reason is simple: Before the FEX is completely operational on Nexus2, it is not possible to configure any host interfaces because they simply do not exist yet!

Example 7-16 shows what happens if you try to configure a host interface in Nexus2 before FEX180 is online. The example also details how you can enable *Fabric Extender preprovisioning* to fix this behavior and effectively provide active-standby parent switches to a Fabric Extender.

Example 7-16 Fabric Extender Preprovisioning

```
! Trying to configure a host interface in FEX180 (which is not operational in Nexus2)
Nexus2(config)# interface ethernet 180/1/1
                                     ^
Invalid range at '^' marker.
! Pre-provisioning FEX180 in Nexus2
Nexus2(config)# slot 180
Nexus2(config-slot)# provision model N2K-C2232P
! Now you can pre-configure a host interface in FEX180
Nexus2(config-slot)# interface ethernet 180/1/1
Nexus2(config-if)# speed 1000
```

Tip The preprovisioning configuration must be done on both switches if you want the host interfaces to be enabled after a switchover from any parent switch to the other.

From a control plane perspective, it might be enough that a single parent switch manages a Fabric Extender while the other waits for a failure. Nevertheless, there are two drawbacks in an active-standby dual-homed topology that challenge its application in real-world scenarios:

- The connection to the standby parent switch is not used for data traffic.

- The transition from one parent switch to the other must wait almost 40 seconds before the Fabric Extender is online.

Figure 7-14 clarifies how virtual PortChannels can overcome both drawbacks and enable active-active dual-homed topologies for Fabric Extenders.

Example 7-17 depicts the vPC configuration in Nexus1. The same configuration was issued in Nexus2, except for the peer keepalive destination IP address that was pointed to 172.19.17.4 (Nexus1 mgmt0 interface).

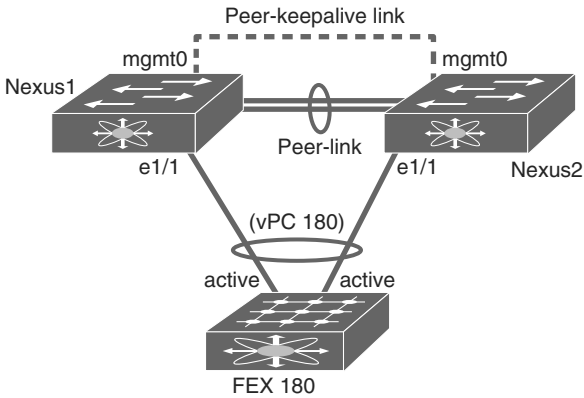


Figure 7-14 Active-Active Dual-Homed Fabric Extender

Example 7-17 Nexus1 vPC Configuration

```

! Observing all vPC configuration parameters in Nexus1
Nexus1# show running-config vpc
[output suppressed]

feature vpc
! Step 1: Configuring vPC domain
vpc domain 1000
  role priority 1
! Step 2: Configuring peer-keepalive link (IP address 172.19.17.4 in Nexus2)
  peer-keepalive destination 172.19.17.5
[output suppressed]
! Step 3: Configuring peer-link (interfaces e1/31 and e1/32)
Nexus1# show running-config interface port-channel 1000, port-channel 180
[output suppressed]
interface port-channel1000
  switchport mode trunk
  spanning-tree port type network
  vpc peer-link
! Step 4: Configuring and reusing PortChannels (interface e1/1 on both switches)
interface port-channel180
  switchport mode fex-fabric
  fex associate 180
  vpc 180

```

In active-active dual-homed scenarios, each Fabric Extender is online for both parent switches, and Examples 7-18 and 7-19 portray this status.

Example 7-18 *Verifying Fabric Extender Status in Nexus1*

```

Nexus1# show fex 180 detail
FEX: 180 Description: FEX0180   state: Online
[output suppressed]
Fabric interface state:
  Po180 - Interface Up. State: Active
  Eth1/1 - Interface Up. State: Active
Fex Port      State  Fabric Port
! Interface e180/1/1 is configured and operational in Nexus1
  Eth180/1/1  Up      Po180
  Eth180/1/2  Down    None
[output suppressed]

```

Example 7-19 *Verifying Fabric Extender Status in Nexus2*

```

Nexus2# show fex 180 detail
FEX: 180 Description: FEX0180   state: Online
[output suppressed]
Fabric interface state:
  Po180 - Interface Up. State: Active
  Eth1/1 - Interface Up. State: Active
Fex Port      State  Fabric Port
! Interface e180/1/1 is configured and operational in Nexus2
  Eth180/1/1  Up      Po180
  Eth180/1/2  Down    None
[output suppressed]

```

In active-active topologies, a parent switch failure does not affect the host interfaces on the Fabric Extender because both vPC peers manage it simultaneously. However, it is a requirement that the FEX configuration (including the host interfaces) is the same on both switches.

Tip You can use the *configuration synchronization* NX-OS feature to replicate selected parts of a configuration to a peer switch (if it is available in your software and hardware combination).

I usually recommend active-active dual-homed topologies for scenarios that require minimum failure effects on server connectivity (such as servers that have only one Ethernet connection, for example).

During a period, it was not possible to deploy vPCs on host interfaces connected to dual-homed active-active Fabric Extenders. Notwithstanding, Enhanced virtual PortChannel (EvPC) capability surpassed that limitation, maintaining a simple configuration principle: Because in active-active dual-homed topologies, the host interfaces are configurable on both parent switches, you only need to deploy a PortChannel for the server interfaces you want to aggregate.

Figure 7-15 depicts the EvPC scenario derived from our topology with the inclusion of another Fabric Extender (FEX190).

Note Consider that the FEX190 configuration on both switches is exactly the same as the one issued for FEX180.

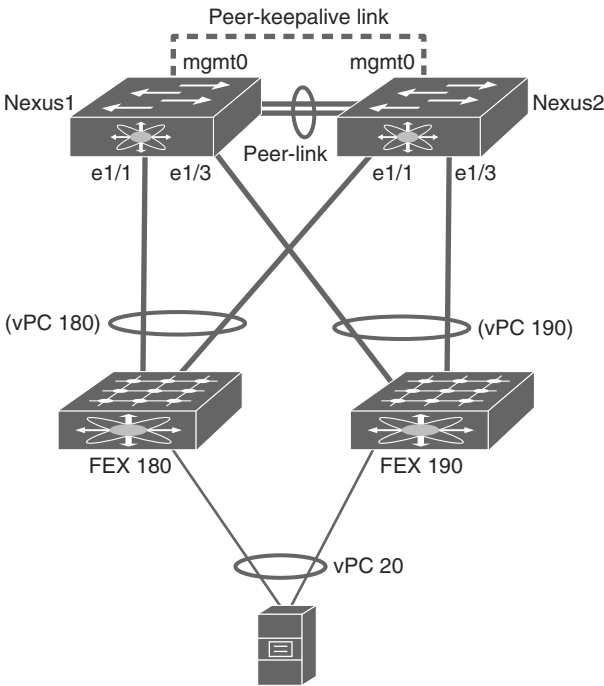


Figure 7-15 Enhanced vPC

The EvPC configuration on Nexus1 is detailed in Example 7-20. As mentioned previously, Nexus2 received the same configuration.

Example 7-20 *Enhanced vPC Configuration in Nexus1*

```

Nexus1# show running-config interface port-channel 20
[output suppressed]
! Configuring the EvPC
interface port-channel20
    switchport access vlan 50

Nexus1# show running-config interface ethernet 180/1/1, ethernet 190/1/1
[output suppressed]
interface Ethernet180/1/1
    switchport access vlan 50
    speed 1000
! Including Ethernet 180/1/1 in the EvPC
    channel-group 20 mode active

interface Ethernet190/1/1
    switchport access vlan 50
    speed 1000
! Including Ethernet 190/1/1 in the EvPC
    channel-group 20 mode active

```

Use Case: Mixed Access Data Center

After your imaginary career achieves the sequential successes described in previous chapters, you are requested to design the server access network of a midrange data center.

This data center is acquiring 160 new rack-mountable servers with two 10 Gigabit Ethernet interfaces (for redundancy), where each server consumes approximately 450 watts. But the design must also support 96 legacy servers that will be decommissioned over the next two years. These old servers have only one Gigabit Ethernet interface and very similar power requirements (although only one-third of the performance).

This company's CIO requires that three principles orient your project: scalability, high availability, and physical optimization. He also presents two boundary conditions:

- Only 7.5 kilowatts can be provided per server cabinet.

- The network cannot have an oversubscription higher than 5:1, even in moments of a connection failure.

Believing that an FEX-based virtualized chassis can be customized to support the specific requirements of most server environments, you present the design depicted in Figure 7-16.




Figure 7-16 *Mixed Access Data Center*

In your design, each server cabinet can accommodate up to 16 servers because these devices consumes together 7.2 kW (leaving 300 W for the Fabric Extender). A pair of Nexus 2232PP Fabric Extenders (with eight 10 Gigabit Ethernet fabric interfaces and 32 10 Gigabit Ethernet host interfaces) can support and optimize the cabling of two server cabinets.

Within these cabinets, each new server is connected to two different Nexus 2232PP Fabric Extenders in a straight-through topology. The fabric interfaces do not use PortChannels to maintain the bandwidth oversubscription of 4:1 if there is a connection failure. Optionally, these servers can leverage vPCs to deploy 20 Gbps of bandwidth with aggregated interfaces.

For the legacy servers, you also consider 16 servers in each cabinet. However, you select one Nexus 2248TP-E Fabric Extender (with four 10 Gigabit Ethernet fabric interfaces and 48 10 Gigabit Ethernet host interfaces) to provide connectivity to three server cabinets and increase high availability through an active-active dual-homed topology. This design assures the network team that, even in the case of a fabric interface failure, the host interface's maximum oversubscription is 4.8:1.

Next, you select a pair of Layer 3 Nexus switches capable of deploying both FEX topologies and acting as the default gateway for the servers. The team is happy to know that only two virtualized switches will be able to manage their entire data center.

During lunch, they reveal that you have just designed their backup data center for the next three years. And later, you are invited to lead their main data center project.

Summary

In this chapter, you have learned that

- Both Top-of-Rack and End-of-Row server access models have advantages and shortcomings. Fabric Extenders leverage the best of both models (cabling optimization and consolidated management, respectively).
- A parent switch connected to one or more Fabric Extenders forms a virtualized access chassis.
- Parent switches and Fabric Extenders use VNTag to deploy virtual interfaces for each FEX host interface.
- FEX host interfaces are RSTP edge ports with BPDU filter and BPDU guard enabled.
- Fabric Extenders can have fabric interface redundancy through static pinning or PortChannels.
- There are two classes of Fabric Extender redundant topologies: straight-through and dual-homed.
- It is possible to provide link aggregation to servers connected to Fabric Extenders in straight-through topologies (standard vPC) and dual-homed topologies (Enhanced vPC).

Figure 7-17 graphically summarizes how Fabric Extenders (connected to a parent switch) can emulate a virtualized modular chassis for a server connection.

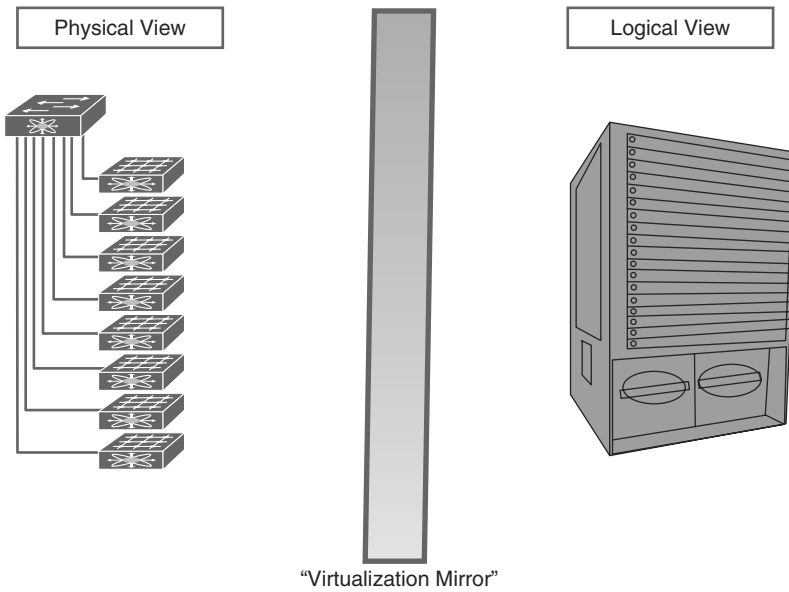


Figure 7-17 *Through the Virtualization Mirror*

Further Reading

Gai, Silvano and DeSanti, Claudio. *I/O Consolidation in the Data Center*. Cisco Press, 2009.

VNTag 101

www.ieee802.org/1/files/public/docs2009/new-pelissier-vntag-seminar-0508.pdf

White Paper: Virtual Machine Networking: Standards and Solutions

www.cisco.com/en/US/prod/collateral/switches/ps9441/ps9902/whitepaper_c11-620065.pdf

IEEE Standards Association

standards.ieee.org/getieee802/download/802.1BR-2012.pdf

This page intentionally left blank

Index

NUMERICS

100GBASE-SR4, 32
10BASE5 (Thicknet), 27
10GBASE2 (Thinnet), 27
10 Gigabit Ethernet, 497
40GBASE-SR4, 32

A

AAA (authentication, authorization, and accounting), 171

abstraction, cloud computing, 786

acceleration, applications, 58, 129, 763-771

accelerators, 111

ACCEPT message, 427

access

 broad network, 794

 core-aggregation-access, 36

 core-distribution-access, 34

 firmware policies, 633

 layers, 35

 lists, 164

See also ACLs

 load balancers, 114

 MAC (Media Access Control), 51

 NUMA (Non-Uniform Memory Access), 566

 ports, 390

 promiscuous access interfaces, 81

 RBAC (Role-Based Access Control), 171

 rules (Cisco ASA 1000V), 761

 servers

models, 287-291

Unified Fabric designs, 536-542

 storage, 391-399

block-based, 392-397

files, 397-398

interfaces, 495

mainframe, 396-397

records, 398-399

 switches, 38

 UCS Manager, 590

 unified access servers

deployment, 509-523

designs, 542-545

 virtual contexts (ACE) management, 171-176

- access-aggregation connections, 38
- access control lists. *See* ACLs
- accounts, user configuration, 175
- ACE (Appliance Control Engine), 185
 - connections, 141-144
 - virtual contexts
 - configuration*, 163-171
 - controlling management access*, 171-176
 - fault tolerance*, 177-178
 - integrating*, 156-161
 - multitenant data centers*, 179-181
 - resource allocation*, 145-156
 - sharing VLANs*, 177
- ACE (Application Control Engine), 109
 - application networking services, 111
 - classification, 110
 - load balancers, 111-134
 - physical connections, 141-145
 - virtual contexts, 139-178
- ACLs (access control lists), 380
 - attributes, 752
 - security policies, creating, 745
 - virtual networks, 692
- activation
 - physical servers, 581
 - servers, provisioning, 583-585
- active-active greenfield data centers, 382-384
- Active Directory. *See* AD
- active nodes, 325
- active-standby dual-homed topologies, 309
- AD (Active Directory), 398
- Adaptable Modular Storage. *See* AMS
- adapters
 - firmware policies, 634
 - network convergence, 509
- Adaptive Security Appliances. *See* ASAs
- Adaptive Security Device Manager. *See* ASDM
- addresses
 - FCIDs (Fibre Channel Identifiers), 429
 - Fibre Channel, 413-415
 - IP (Internet Protocol)
 - assigning*, 59
 - Layer 4/Layer 7 switches*, 120
 - management*, 610
 - virtualization clusters*, 668
 - VSG (Virtual Security Gateway)*, 752
 - LBAs (Logical Block Addresses), 394
 - load balancer translation, 124-127
 - LUNs (logical unit numbers), 393
 - MAC (Media Access Control), 51, 581
 - BPDU frames*, 65
 - FabricPath*, 331
 - flooding*, 340
 - FPMA (Fabric-Provided MAC Address)*, 508
 - learned in MPLS*, 346
 - OTV (Overlay Transport Virtualization)*, 355
 - service profiles*, 587
 - sharing resources*, 85
 - SPMA (Server-Provided MAC Address)*, 508
 - tables*, 691
 - NAT (Network Address Translation), 466
 - overlapping, 87-89
 - pWWN (port world wide name), 465

- remapping, 17
- URLs (Uniform Resource Locators), 116
- VIP (virtual IP), 116
- Address Resolution Protocol. *See* ARP**
- adjacency, OTV (Overlay Transport Virtualization), 370**
- adjacent pairs, 28**
- Admin tab (UCS Manager), 591**
- Advanced Encryption Standard (AES), 133**
- Advanced VPLS. *See* A-VPLS**
- advantages, 21**
 - ACE virtual contexts, 110
 - Fabric Extenders, 288
 - load-balancer deployments, 114
 - network services, 736
 - service profiles, 582
 - virtualization, 454, 494
 - VLANs/VRF, 45
 - VM-FEX (Virtual Machine Fabric Extender), 657
 - vPCs (virtual PortChannels), 231
 - VSANs (virtual storage-area networks), 409
- AEDs (authoritative edge devices), 365**
- AES (Advanced Encryption Standard), 133, 328**
- aggregation**
 - failures, 38
 - layers, 36
 - links, 234-240
 - MC-LAG (Multi-Chassis Link Aggregation Group), 351
 - virtual networking, 684-688
- algorithms**
 - acceleration, 763
 - slow-start, 116
- allocation**
 - interfaces, 190
 - memory, 210
 - resources, 17
 - VDCs (*virtual device contexts*), 202-211
 - virtual contexts (ACE)*, 145-156
 - VLANs (*virtual LANs*), 206
 - VRF (*Virtual Routing and Forwarding*), 101-103
 - virtual context creation, 145
- Amazon Web Service. *See* AWS**
- American National Standards Institute. *See* ANSI**
- American National Standards Institute and Telecommunications Industry Association (ANSI/TIA), 29**
- AMS (Adaptable Modular Storage), 792**
- analysis, traffic, 58**
- ANM (Application Network Manager), 147**
- ANSI (American National Standards Institute), 410**
- ANSI/TIA (American National Standards Institute and Telecommunications Industry Association), 29**
- ANSI/TIA-942 standards, 40-41, 287**
- answer files, 585**
- Any Configuration option, 595**
- Any Transport over MPLS. *See* AToM**
- API (application programming interface), 661**
- APIC (Asynchronous Inter-Process Communication), 662**
- appliances**
 - ACE (Appliance Control Engine) connections, 141-144
 - ASAs (Adaptive Security

- Appliances), 179
- Cisco ASA 1000v, 754-762
- network services, 738
 - See also* networks; services
- VSG (Virtual Security Gateway), 742-763
- Application Control Engine.** *See* ACE
- application control software,** 142
- Application Network Manager.** *See* ANM
- application programming interface.** *See* API
- applications,** 19
 - acceleration, 129, 763-771
 - batch processing, 321
 - delivery controllers, 765
 - environment independence, 138-139
 - load balancers, 127-130
 - networking services, 111
 - resilience, 37
 - scaling, 113
 - virtual LUNs, 405
- application-specific integrated circuits.** *See* ASICs
- applying load balancers,** 111-134
- AppNav,** 765
- arbitrated loops,** 412
- architecture**
 - Cisco ASA 1000v, 755
 - Cisco Nexus 1000v, 661-663
 - control planes, 742
 - data centers, 5-7
 - Fabric Extender, 292
 - layers, 35-36
 - NX-OS, 189
 - See also* NX-OS switches
 - SCSI (Small Computer Systems Interface), 395
 - servers, 560-562
 - services. *See* services
 - SSA (Serial Storage Architecture), 396
 - Storage Reference Architecture, 792
 - virtualization, extending, 185
 - VMDC (Virtualized Multiservice Data Center), 792
 - WAAS (Wide Area Application Services), 764
- areas,** 20, 21
 - operational, 5-7
 - security, 47
 - technologies, 18-21
 - ACE virtual contexts,* 110
 - Fabric Extenders,* 288
 - service profiles,* 582
 - virtualization,* 454, 494
 - VM-FEX (Virtual Machine Fabric Extender),* 657
 - vPCs (virtual PortChannels),* 231
 - VSANs (virtual storage-area networks),* 409
- VLANs/VRF, 45
- ARP (Address Resolution Protocol),** 49, 701
- arrays**
 - clusters, 403
 - disk, 389-390
 - zones, 439
- ASAs (Adaptive Security Appliances),** 179, 735, 737, 754-762
- ASDM (Adaptive Security Device Manager),** 754
- ASICs (application-specific integrated circuits),** 103
- assignment**
 - HBAs (host bus adapters), 595
 - interfaces to Fibre Channel, 413
 - IP addresses, 59
 - servers, 606-607

UUID (Universally Unique Identifier), 594
 values, default path cost for switches, 64
 VLAN (virtual local-area network) interfaces, 52
association
 service profiles to servers, 612-619
 VRF (Virtual Routing and Forwarding), 90
asymmetric connection management, 122
Asynchronous Inter-Process Communication (AIPC), 662
asynchronous replication, 402
ATA Packet Interface (ATAPI), 397
ATAPI (ATA Packet Interface), 397
Atlas Team, 9
AToM (Any Transport over MPLS), 333
attachments, interfaces, 343
attenuation, 29
attributes
 ACLs (access control lists), 752
 VMs (virtual machines), 751
authentication
 offloading servers, 130
 UCS Manager, 589
authentication, authorization, and accounting. See AAA
authoritative edge devices. See AEDs
automation
 cloud computing, 789-792
 provisioning, 581
autostate (VLANs), 60
auxiliary memory, 8, 388
availability, 25
 distributed data centers, 321
 geoclusters, 323
 networks, 37

 servers, 559
avoidance, loops, 326, 330
 Cisco Nexus 1000v switches, 685
 OTV (Overlay Transport Virtualization), 365-373
A-VPLS (Advanced VPLS), 350
AWS (Amazon Web Service), 794

B

backbone cabling, 41
backup interfaces, 495
bandwidth
 ETS (Enhanced Transmission Selection), 500
 load balancer performance, 135
basic input/output system. See BIOS
batch processing applications, 321
BDs (bridge domains), 208
behavior
 broadcast frames, 61
 disabling, 164
 hubs, 48
 paths, 47
 private VLAN ports, 79
 server responses, 118
 switches, 48
 virtual contexts (ACE), 150
 VLAN (virtual local-area network) switches, 50
Bell, Alexander Graham, 28
benefits, 25
 of networks, 42-44
 of virtualization, 2
best practices
 deployment, 36
 predictors, selecting, 117
BF (Build Fabric) frame, 423
BGP (Border Gateway Protocol), 87, 334

BI (Business Intelligence), 3

binding

labels, 337

MPLS (Multiprotocol Label Switching) connections, 353

VNMC (Virtual Network Manager Center), 760

biological taxonomy, 21

See also taxonomies

BIOS (basic input/output system), 581

configuration, 627-633

firmware policies, 635

server domains, 585

updating, 636

x86 architecture, 563

blade servers, 488, 540-542

blade switches, 474

block-based access, 392-397

Blocked state (STP), 70

blocks, 391

aggregation, 36

BPDU (bridge protocol data units), 684

LBAs (Logical Block Addresses), 394

SMB (Server Message Block), 398

Boggs, Dave, 26

booting

configuration, 225

order, servers, 604-605

BOOTP (Bootstrap Protocol), 51

Bootstrap Protocol. *See* BOOTP

Border Gateway Protocol. *See* BGP

Boucher, Larry, 392

BPDU (bridge protocol data units), 65, 233

blocks, 684

Fabric Extender, 300

Layer 2 extensions, 329

bridge domains. *See* BDs

Bridge-group Virtual Interfaces. *See* BVIs

bridge protocol data units. *See* BPDUs

bridges

DCB (Data Center Bridging), 497-503

Ethernets, 47

IDs, 63

root, 62

virtual contexts (ACE) design, 158-160

Broadcast Alias service, 421

broadcasts

ARP (Address Resolution Protocol), 49

domain extensions, 325

Ethernet loops, 62

gatekeepers, 365

vPCs (virtual PortChannels), 243

broad network access, 794

buffers

buffer-to-buffer credits calculation, 418, 419

PAUSE frames, 499

Build Fabric (BF) frame, 423

building blocks, Fibre Channel, 416

buses

Ethernets, 47

SCSI (Small Computer Systems Interface), 393

shared, 392

x86

architecture, 563

server expansion, 569-571

Business Intelligence. *See* BI

Bus&Tag communication, 396

BVIs (Bridge-group Virtual Interfaces), 145, 158

C

cables, 493

See also connections

backbone cabling, 41

coaxial (Ethernet protocol), 27-28

convergence, 495-497

DCB (Data Center Bridging), 497-503

direct-attach twinaxial (Ethernet protocol), 32-33

FCoE (Fibre Channel over Ethernet), 504-509

horizontal, 289

intra-rack cabling, 40

optical fiber (Ethernet protocol), 29-32

optimization, 44

trends, 34

twisted-pair, 28-29, 34

cache, 129, 390

memory, 9

WCCP (Web Cache Control Protocol), 765

calculations

sequence numbers, 133

sliding windows, 133

Canonical Format Indicator. *See* CFI

carrier sense multiple access collision detect. *See* CSMA/CD

Catalyst 3750 switches, 240

Catalyst 6500 switches

ACE module connections, 144

virtual contexts (ACE), 135

Virtual Switch System, 240

categories, ANSI/TIA-568 twisted-pair, 30

CCP (Cisco Cloud Portal), 797

CDBs (command descriptor blocks), 393

CDP (Cisco Discovery Protocol), 601

CE (customer edge) devices, 338

CEE (Converged Enhanced Ethernet), 497

central processing units. *See* CPUs

CEOs (chief executive officers), 6

Certification no. 10349, 227

certifications, governance, 583

CFI (Canonical Format Indicator), 53

CFOs (chief financial officers), 6

CFS (Cisco Fabric Services), 243, 466

changeto command, 147

channels

Fibre Channel, 410

See also Fibre Channel

I/O (input/output), 388

chassis management circuit. *See* CMC

checking status of NPV (N_Port Virtualization), 477

checksums, 133

chief executive officers (CEOs), 6

chief financial officers (CFOs), 6

chief information officers (CIOs), 2, 6

chief security officers (CSOs), 6

chipsets, x86 architecture, 563

CIAC (Cisco Intelligent Automation for Cloud), 797

CIFS (Common Internet File System), 398, 763

CIMC (Cisco Integrated Management Controller), 588

firmware policies, 633

versions, 634

CIOs (chief information officers), 2, 6

CISC (Complex Instruction Set Computing), 561

Cisco ASA 1000V switches, 754-762

access rules, 761

configuration, 757

- installation, 755-757
 - security policy configuration, 761
 - traffic, sending to, 758
 - VNMC (Virtual Network Manager Center) registration, 756
- Cisco Cloud Portal (CCP), 797**
- Cisco Discovery Protocol. *See* CDP**
- Cisco Fabric Services. *See* CFS**
- Cisco Integrated Management Controller. *See* CIMC**
- Cisco Intelligent Automation for Cloud (CIAC), 797**
- Cisco Nexus 1000v switches**
 - architecture, 661-663
 - communication modes, 663-664
 - deployment, 666-683
 - loop avoidance, 685
 - online migrations, 693-697
 - port profiles, 664-666
- Cisco Open Network Environment. *See* ONE**
- Cisco Process Orchestrator (CPO), 798**
- Cisco Server Provisioner (CSP), 798**
- cladding, 29**
- classes**
 - maps
 - Layer 7*, 167
 - VIP (virtual IP)*, 169
 - resources, 150, 155
 - of service (Fibre Channel), 420
- classification**
 - ACE (Application Control Engine), 110
 - Ethernet ports, 72
 - Fabric Extenders, 288
 - Layer 2 extensions, 320
 - network services, 736
 - service profiles, 582
 - VDCs (virtual device contexts), 183
 - virtualization, 14-21, 454, 494
 - VLANs (virtual local-area networks), 45
 - VM-FEX (Virtual Machine Fabric Extender), 657
 - vPCs (virtual PortChannels), 231
 - VRF (Virtual Routing and Forwarding), 45
 - VSANs (virtual storage-area networks), 409
- CLI (command-line interface) 20, 142**
 - MPC (Modular Policy CLI), 162
 - prompts, 198
 - VDCs (virtual device contexts), 192
- client-by-client configuration, 129**
- clients, 111**
- clock generators, x86 architecture, 563**
- cloning, 638-639**
- cloud computing, 785-786**
 - automation, 789-792
 - deployment, 799-800
 - implementation, 797-798
 - networks, 800, 802-804
 - ONE (Cisco Open Network Environment), 804
 - OpenStack, 801-802
 - overview of, 793-797
 - SDNs (Software-Defined Networks), 800-801
 - standardization, 789-792
 - virtual data centers, 786-789
- Cloud Services Routers. *See* CSRs**
- clusters**
 - arrays, 403
 - interfaces, 700
 - IP (Internet Protocol) addresses, 668
 - security policies, 737
 - servers, 324
 - virtualization, 737

- CMC (chassis management circuit), 588
- CMS (Conversational Monitor System), 10
- CNAs (converged network adapters), 585
- CNMs (congestion notification messages), 503
- coarse wavelength-division multiplexing (CWDM), 328
- coaxial cables (Ethernet protocol), 27-28
- COBIT (Control Objectives for Information and Related Technologies), 583
- code, SCSI commands, 393
- cold-standby sites, 322
- collisions
 - domains, 48
 - VLANs (virtual local-area networks), 377-379
- command descriptor blocks. *See* CDBs
- command-line interfaces. *See* CLIs
- commands
 - changeto, 147
 - fex associate, 297
 - Format Unit, 393
 - Inquiry, 393
 - lacp, 237
 - mpls ip, 335
 - Nexus 7000, 61
 - Read, 393
 - Read Capacity, 393
 - reload vdc, 215
 - SCSI (Small Computer Systems Interface), 393
 - Send Diagnostic, 393
 - show interface, 459
 - switchmode fex-fabric, 297
 - Test Unit Ready, 393
 - vlan all, 55
 - vrf member interface, 91
 - Write, 393
- Common Internet File System (CIFS), 398
- Common Spanning Tree. *See* CST
- communication
 - errors, 48
 - Ethernets, 58
 - modes (Cisco Nexus 1000v switches), 663-664
 - UDP (User Datagram Protocol), 120
- Complex Instruction Set Computing. *See* CISC
- components
 - of disk arrays, 389
 - FCIP (Fibre Channel over IP), 455
- compression
 - HTTP (Hypertext Transfer Protocol), 130, 134
 - offloading servers, 130
 - PLZ (Persistent Lempel-Ziv), 764
- computer rooms (ANSI/TIA-942 standards), 40
- concurrent connections, 135
- confidentiality, 326
- configuration
 - BIOS (basic input/output system), 581, 585, 627-633
 - boot, 225, 604
 - BVIs (Bridge-group Virtual Interfaces), 158
 - Cisco ASA 1000V, 757
 - client-by-client, 129
 - complexity, 137
 - CSRs (Cloud Services Routers), 773
 - DHCP (Dynamic Host Configuration Protocol), 51
 - domains, 175
 - EEM (Event Embedded Manager), 349

- EIGRP (Enhanced Interior Gateway Routing Protocol), 96
- EoMPLS (Ethernet over MPLS), 338
- Ethernet 10/10, 197-198
- EvPC (Enhanced virtual PortChannel), 315
- Fabric Interconnect, 588
- FabricPath, 272
- FCIP (Fibre Channel over IP), 457-460
- Flex-Attach, 487
- F-Trunks, 480
- interfaces
 - out-of-band management switches, 143*
 - production switches, 142*
- IPMI (Intelligent Platform Management Interface), 609
- LACP (Link Aggregation Control Protocol), 237
- load balancers
 - policies, 132*
 - servers, 115*
- MPLS (Multiprotocol Label Switching), 335
- MST (Multiple Spanning Tree), 201
- multihop FCoE (Fibre Channel over Ethernet), 523-535
- nondisruptive reconfiguration of fabric, 423
- NPIV (N_Port ID Virtualization), 476-482
- operational policies, 608
- OSPF (Open Shortest Path First), 94
- OTV (Overlay Transport Virtualization), 359-364
- PortChannels, 303, 306
 - Layer 3, 236*
 - STP (Spanning Tree Protocol), 236*
- port profiles, 664
- predictors, 116
- private VLANs, 80-82
- probes, 164-167
- promiscuous access interfaces, 81
- Q-in-Q tunnel interfaces, 378
- RAID (Redundant Array of Independent Disks), 399-401
- security policies (Cisco ASA 1000V), 761
- server farms, 166
- SPAN sessions, 205-206
- SVIs (Switch Virtual Interfaces), 59
- Switch1, 196-197
- trunk interfaces, 55
- unified access servers
 - single-context switches, 510-518*
 - storage VDCs, 519-523*
- users
 - accounts, 175*
 - login, 176*
- VDCs (virtual device contexts), 190-202
- VFC (Virtual Fibre Channel), PortChannels, 528-531
- vHBAs (virtual HBAs), 596
- virtual contexts (ACE), 140, 163-171
- virtual networking, 658-661
- VLANs (virtual local-area networks), 85, 209
- vNICs (virtual NICs), 599
- vPath
 - Nexus 1000v switches, 748, 758*
 - WAAS (Wide Area Application Services), 766*
- vPCs (virtual PortChannels), 247-254
- VRF (Virtual Routing and Forwarding), 90-91, 99, 199
- VSG (Virtual Security Gateway) fire-

- walls, 746
- VSS (Virtual Switch System), 241
- vSwitch (Virtual Switch), 673
- VXLANs (Virtual eXtensible Local Area Networks), 703
- vZones (virtual zones), 754
- congestion**
 - control, 134
 - QCN (Quantized Congestion Notification), 503
- congestion notification messages.** *See* CNMs
- connected data centers, internal routing in, 380-382**
- connections, 18**
 - access-aggregation, 38
 - ACEs (Application Control Engines), 141-145
 - concurrent, 135
 - cross-connections, 41
 - dedicated, 392
 - desktops, 34
 - devices, 32
 - EOBC (Ethernet Out-of-Band Channel), 144
 - EoR (End-of-Row) designs, 289-291
 - external, virtual networking, 684-688
 - Fabric Extenders, 296-299
 - hosts, 35
 - internal routing, Layer 2 extensions, 380-382
 - interswitch, 234
 - IP (Internet Protocol), 200
 - See also* IP
 - Layer 2 interfaces, 343
 - least-connections predictors, 116
 - limitations, 118
 - management, 122-124
 - migration, 289
 - MPLS (Multiprotocol Label Switching), 353
 - multiple (STP), 233
 - multisite interconnections (FabricPath), 331
 - N7K switches, 524
 - Nexus 1000V switches, 741
 - NIC (network interface card) teaming, 239
 - NPV (N_Port Virtualization), 476
 - See also* NPV
 - offloading servers, 130
 - optical, Ethernet extensions over, 327-332
 - optimization,
 - peer keepalive links, 248
 - physical, 495
 - ports
 - switches, 69*
 - VDCs (virtual device contexts), 192*
 - pseudowires, 343
 - servers, 34
 - state communications, 325
 - TCP (Transmission Control Protocol), 120
 - termination, 133
 - Top-of-Rack (ToR) designs, 289-291
 - UDP (User Datagram Protocol), 120
 - vCons (virtual network interface connections), 602
 - vPath, 768
- connections per second (cps), 135**
- connectors**
 - RJ-45, 29
 - twinax cables, 32
- consolidation, 42**
 - management, 44
 - SANs (storage area networks), 447-450

- switches, 187
- VM-FEX (Virtual Machine Fabric Extender), 705
- content, security, 58**
- contention**
 - cloud computing, 803
 - PFC (priority-based flow control), 498
- Content Services Switch. *See* CSS**
- Content Switching Module. *See* CSM**
- contexts, virtual (ACEs), 109**
 - See also* ACEs
- control architecture (Cisco ASA 1000v), 755**
- controllers, 390**
 - application delivery, 765
 - CIMC (Cisco Integrated Management Controller), 588
 - fabric, 421
 - FCoE (Fibre Channel over Ethernet), 505
 - memory (x86 architecture), 562
 - SDNs (Software-Defined Networks), 800
- control planes, 20**
 - architecture, 742
 - FabricPath, 269-272
- Control Program (CP), 10**
- control units. *See* CUs**
- converged access models, 542**
- converged aggregation models, 543**
- Converged Enhanced Ethernet. *See* CEE**
- converged network adapters. *See* CNAs**
- convergence**
 - I/O (input/output), 495-497
 - network adapters, 509
 - RSTP (Rapid Spanning Tree Protocol), 72
 - speed, 232
 - STP (Spanning Tree Protocol), 66
- Conversational Monitor System (CMS), 10**
- core-aggregation-access, 36**
- core-distribution-access, 34**
- cores, 29, 564**
- costs, default path cost for switches, 64**
- counters, sequences, 415**
- CP (Control Program), 10**
- CPO (Cisco Process Orchestrator), 798**
- cps (connections per second), 135**
- CPUs (central processing units), 4, 388**
 - utilization, 203
 - virtual memory, 9
 - x86 architecture, 562
 - x86 servers, 564-566
- CRC (cyclic redundancy check), 417**
- Create Host Firmware Package wizard, 634**
- Create Service Profile (expert) wizard, 594, 626, 637**
- creating. *See* formatting**
- credits, buffer-to-buffer, 418-419**
- CRM (Customer Relationship Management), 3**
- cross-connections, 41**
- cross-switch**
 - aggregation techniques, 342
 - PortChannels, 240-241
- crosstalk interference, 28**
- CSM (Content Switching Module), 139**
- CSMA/CD (carrier sense multiple access collision detect), 26**
- CSOs (chief security officers), 6**
- CSP (Cisco Server Provisioner), 798**
- CSRs (Cloud Services Routers), 735**
 - 1000V connectivity parameters, 771

configuration, 773
CSS (Content Services Switch), 139
CST (Common Spanning Tree), 75
CUs (control units), 396
customer edge. *See* CE
Customer Relationship Management.
See CRM
**CWDM (coarse wavelength-division
multiplexing)**, 328
cyclic redundancy check. *See* CRC

D

**database management system
(DBMS)**, 398
Data Center Bridging. *See* DCB
**Data Center Bridging eXchange pro-
tocol**. *See* DCBX protocol
Data Center Ethernet. *See* DCE
Data Center Interconnect. *See* DCI
data centers
architecture, 5-7
blade servers, 488
convergence, 495-497
definitions, 2-7
evolution of, 3-5
internal routing (Layer 2 extensions),
380-382
limitations of, 5
load balancers, proliferation in, 135-
139
merging, 731-733
mixed access (Fabric Extender), 315-
317
multitenant, virtual contexts (ACE),
179-181
optimization, 231
origins of virtualization, 8-13
segmentation, 103-104
storage devices, 387-391
VDCs (virtual device contexts), 225-
227
virtual, 781-783, 785-789
See also cloud computing
virtual contexts (ACE), integrating,
156-161
data confidentiality, 326
Data Encryption Standard (DES),
133
data paths, 738-740
See also vPath
data planes, 20, 266-269
data rates (Ethernet standards), 33
Data Redundancy Elimination (DRE),
764
Data Warehouses. *See* DWs
**DBMS (database management sys-
tem)**, 398
DCB (Data Center Bridging), 493,
497-503
**DCBX (Data Center Bridging
eXchange) protocol**, 501-503
**ETS (Enhanced Transmission
Selection)**, 500-501
PFC (priority-based flow control),
498-499
**QCN (Quantized Congestion
Notification)**, 503
**DCBX (Data Center Bridging
eXchange) protocol**, 501-503, 518
DCE (Data Center Ethernet), 497
**DCE (Distributed Computing
Environment)**, 594
DCI (Data Center Interconnect), 324
STP (Spanning Tree Protocol), 326
switches, 328
DCNM-SAN topologies, 531
DDR (Double Data Rate), 567
DEC (Digital Equipment Company),
26
DEC, Intel, Xerox (DIX), 26

dedicated connections, 392

dedicated processors (x86 architecture), 563

default domains, 175

default gateways, 11

default HA policies, 217

default path cost for switches, 64

default VDC resource allocation, 212

definitions

data centers, 2-7

dynamic vNIC connection policies, 708

Fibre Channel, 410-420

network services, 599-608, 735-738

service profiles, 595-598

storage, 8

taxonomies, 14

vHBAs (virtual HBAs), 596

virtualization, 1, 8, 12

VLANs (virtual local-area networks), 49-56

vPCs (virtual PortChannels), 242-248

VRF (Virtual Routing and Forwarding), 90-91

VSANs (virtual storage-area networks), 430-447

delivery controllers, applications, 765

demilitarized zones. *See* DMZs

dense wavelength-division multiplexing (DWDM), 328

deployment

Cisco Nexus 1000v switches, 666-683

clients, 111

cloud computing, 796, 799-800

DNS load balancers, 112

ISL (Inter-Switch-Link), 56

redundant VPLS (Virtual Private LAN Service), 348

routing, 58

services, 111

switches (STP versions), 62

topologies, 36

UCS (Unified Computing System), 578

unified access servers, 509-523

VLANs (virtual local-area networks), 52

VM-FEX (Virtual Machine Fabric Extender), 706-720

vSphere (VMware), 709-710

VXLANs (Virtual eXtensible Local Area Networks), 699

WWNs (World Wide Names), 486-488

Deploy OVF Template vCenter wizard, 745

DES (Data Encryption Standard), 133

DeSanti, Claudio, 509

design, 26

See also configuration

EoR (End-of-Row), 39, 44

OTV (Overlay Transport Virtualization), 373-377

partitioning, 47

reverse proxies, 128-130

router-on-a-stick, 55, 58

STP (Spanning Tree Protocol), 232

three-layer campus, 34

ToR. *See* ToR (Top-of-Rack) designs

ToR (Top-of-Rack), 39, 44

Unified Fabrics, 535-545

virtual contexts (ACE)

bridges, 158-160

one-armed, 160-161

routers, 156-158

VLANs (virtual local-area networks), 120

Designated Ports Selection phase, 68

desktop connections, 34

detection

loops, 193

volume, 440

devices

See also hardware

ASDM (Adaptive Security Device Manager), 754

connections, 32

edge, 357

groups, 47

load balancer concepts, 115-120

PODs (pool of devices), 281-284

server farms for websites, 168

storage, 4, 18, 387-391

See also storage

DF (Don't Fragment) bit sets, 356

DHCP (Dynamic Host Configuration Protocol), 51, 585, 762

DIA (Domain ID Identifier) frames, 423

different security zones. *See* DMZs

Digital Equipment Company (DEC), 26

DIMMs (Dual In-Line Memory Modules), 568

direct-attach twinaxial cables, 32-33

Direct Server Return, 123-124

Disabled state (STP), 70

disabling

behaviors, 164

legacy extension, 371

disaster recovery plans, 321

discovery (Fabric Extenders), 297

disk arrays, 389-390, 439

disk enclosures, 390

dissimulation, 12

Distributed Computing Environment.
See DCE

distributed data centers, 319, 321

See also data centers

1970s-1980s, 321-322

1990s-mid-2000s, 322-324

mid-2000s to present, 324

Distributed Virtual Switch. *See* DVS

distribution

core-distribution-access, 34

layers, 35

power, 7

DIX (DEC, Intel, Xerox), 26

DL (Don't Learn), 268

DMZs (demilitarized zones), 136,
225-226, 737

DNS (Domain Name System), 111

Domain ID Identifier (DIA) frames,
423

Domain Name System. *See* DNS

domains

BDs (bridge domains), 208

broadcast extensions, 325

collisions, 48

creating, 175

failures, 47

ID distribution, 438

infrastructure, 585, 594

servers, 584-595

UCS (Unified Computing System),
586

vPCs (virtual PortChannels), 248

Don't Fragment (DF) bit sets, 356

Don't Learn. *See* DL

Double Data Rate. *See* DDR

DRAM (dynamic RAM), 566

DRE (Data Redundancy Elimination),
764

dual-homed fabric extenders, 537

dual-homed topologies (Fabric
Extender), 309-315

Dual In-Line Memory Modules. *See*

DIMMs
 Dual NAT, 125-126
 DVS (Distributed Virtual Switch), 661
 DWDM (dense wavelength-division multiplexing), 328, 455
 DWs (Data Warehouses), 3
 Dynamic Host Configuration Protocol. *See* DHCP
 dynamic interfaces, provisioning, 664-666
 dynamic RAM. *See* DRAM
 Dynamic Resource Scheduling, 669
 dynamic vNICs, enabling, 707-709

E

ECC (Error Correction Code), 569
 ECMP (Equal Cost Multipath), 803
 e-commerce, 3
 EDA (Equipment Distribution Area), 41
 edge devices, 357, 754-762
 edge-ports, 72
 Edit Virtual Machine Settings option, 690
 EEM (Event Embedded Manager), 225, 349
 efficiency, 25
 EFP (Exchange Fabric Parameters) phase, 423
 Egress Tunnel Router (ETR), 779
 EIA (Electronic Industries Alliance), 287
 EID (Endpoint Identifier), 779
 EIGRP (Enhanced Interior Gateway Routing Protocol), 87, 96
 EISL (Enhanced Inter-Switch Link), 434
 elasticity, 794
 electromagnetic interfaces. *See* EMIs
 Electronic Industries Alliance (EIA), 287
 elements
 FCoE (Fibre Channel over Ethernet), 505-506
 HDDs (hard disk drives), 389
 load balancer configuration, 115
 ELP (Exchange Link Parameters) phase, 423
 ELS (Extended Link Service) frames, 422
 email filters, 129
 EMIs (electromagnetic interfaces), 28
 emulation, 12, 15, 21
 ACE virtual contexts, 110
 Fabric Extenders, 288
 LAN Emulation (LANE), 803
 network services, 736
 service profiles, 582
 virtualization, 454, 494
 VLANs/VRF, 45
 VM-FEX (Virtual Machine Fabric Extender), 657
 VMs (virtual machines), 659
 vPCs (virtual PortChannels), 231
 VSANs (virtual storage-area networks), 409
 enabling
 acceleration, 769
 LACP processes, 237
 NPV (N_Port Virtualization), 533
 vNICs (virtual NICs), 707-709
 Encapsulated Remote Switch Port Analyzer. *See* ERSPAN
 encapsulation
 frames, 739
 GRE (generic routing encapsulation), 352-354
 ISL (Inter-Switch-Link), 56
 OTV (Overlay Transport Virtualization), 356

- Q-in-Q, 378
- enclosures, disk, 390
- encryption
 - HTTPS (Secure HTTP), 130
 - offloading servers, 130
 - standards, 133
- End-Host mode, 592
- end of frame (EOF), 417
- End-of-Row (EoR) designs. *See* EoR designs
- Endpoint Identifier (EID), 779
- end-to-end congestion notification, 503
- end-to-end SSL (Secure Sockets Layer), 132
- Enhanced Interior Gateway Routing Protocol. *See* EIGRP
- Enhanced Inter-Switch Link. *See* EISL
- Enhanced Quad Small Form Factor Pluggable (QSFP+), 32
- Enhanced Small Form Factor Pluggable (SFP+), 32
- Enhanced Transmission Selection. *See* ETS
- Enhanced virtual PortChannel (EvPC), 314
- ENode, 505
- Enterprise Resource Planning. *See* ERP
- Enterprise System Connection (ESCON), 396
- entrance rooms (ANSI/TIA-942 standards), 40
- environment independence, applications, 138-139
- EOBC (Ethernet Out-of-Band Channel), 144
- EOF (end of frame), 417
- EoMPLS (Ethernet over MPLS), 319, 333, 338-342
- EoR (End-of-Row) designs, 39, 44, 289-291
- EPP (Exchange Peer Parameter)
 - frame, 423
- Equal Cost Multipath (ECMP), 803
- Equipment Distribution Area (EDA), 41
- Equipment tab (UCS Manager), 591
- ERP (Enterprise Resource Planning), 3
- Error Correction Code. *See* ECC
- errors, communication, 48
- ERSPAN (Encapsulated Remote Switch Port Analyzer), 203
- ESCON (Enterprise System Connection), 396
- ESC (Exchange Switch Capabilities)
 - phase, 423
- estimation, memory, 211
- Etheranalyzer sessions, 225
- EtherChannel, 43
- Ethernet 10/10 configuration, 197-198
- Ethernet Adapter Policy option, 601
- Ethernet Out-of-Band Channel. *See* EOBC
- Ethernet over MPLS. *See* EoMPLS
- Ethernets
 - 10 Gigabit, 497
 - bridges, 47
 - broadcast loops, 62
 - CEE (Converged Enhanced Ethernet), 497
 - communication, 58
 - extensions
 - IP (Internet Protocol)*, 352-377
 - over optical connections*, 327-332
 - FCoE (Fibre Channel over Ethernet), 504-509
 - frames, forwarding, 267
 - interfaces, 591
 - LACP modes*, 237

- VSM (Virtual Supervisor Module)*, 662
 - Layer 2. *See* Layer 2
 - MAC (Media Access Control), 505
 - port classification, 72
 - protocols, 25-26
 - coaxial cables*, 27-28
 - direct-attach twinaxial cables*, 32-33
 - media*, 27
 - optical fiber*, 29-32
 - overview of*, 26-34
 - timelines*, 33-34
 - twisted-pair cables*, 28-29
 - segments, 47
 - switches, 48, 184
 - See also* VDCs
 - VEM (Virtual Ethernet Module), 739, 740
 - ETR (Egress Tunnel Router), 779
 - ETS (Enhanced Transmission Selection), 500-501
 - Event Embedded Manager. *See* EEM
 - evolution
 - of data centers, 3-5
 - of hardware (x86 servers), 562-572
 - of networks, 25, 281-284
 - of servers, 559-560
 - EvPC (Enhanced virtual PortChannel), 314
 - Exchange Fabric Parameters (EFP)
 - phase, 423
 - Exchange Link Parameters (ELP)
 - phase, 423
 - Exchange Peer Parameter (EPP)
 - frame, 423
 - exchanges (Fibre Channel), 415-417
 - Exchange Switch Capabilities (ESC)
 - phase, 423
 - EXEC prompt, 147
 - expansion busses (x86 servers), 569-571
 - exporting plug-ins, 712
 - Extended Link Service (ELS) frames, 422
 - extending virtualization, 184-186
 - extensibility, 26
 - Extensible Markup Language. *See* XML
 - extensions
 - broadcast domains, 325
 - Ethernet over optical connections, 327-332
 - FCoE (Fibre Channel over Ethernet), 545
 - full-mesh, 329
 - Layer 2, 319
 - FabricPath*, 330-332
 - internal routing*, 380-382
 - IP (Internet Protocol)*, 352-377
 - MPLS (Multiprotocol Label Switching)*, 332-351
 - OTV (Overlay Transport Virtualization)*, 354-377
 - overview of*, 324-327
 - solutions*, 382-384
 - VLAN (virtual local-area network) identifiers*, 377-379
 - vPCs (virtual PortChannels)*, 328-330
 - VPLS (Virtual Private LAN Service)*, 342-351
 - networks, 44
 - external connections, 684-688
 - external security zones, 225
- ## F
-
- fabric, 412
 - Fibre Channel, 420-429

- initialization, 422-424
- Unified Fabrics design, 535-545
- United Fabric, 495
- Fabric Controller service, 421**
- Fabric Extender.** *See* FEX
- Fabric Failover, 600**
- Fabric Interconnect, 587-588**
- Fabric Interface (FIF), 292**
- Fabric Login, 437, 476**
- failures**
 - See* FabricPath
 - control planes, 269-272
 - data planes, 266-269
 - Layer 2 extensions, 330-332
 - multipathing with, 43
 - STP (Spanning Tree Protocol), 272-276, 330
- Fabric-Provided MAC Address.** *See* FPMA
- Fabric Shortest Path First.** *See* FSPF
- failover (Fabric Failover), 600**
- failures**
 - See also* troubleshooting
 - aggregation, 38
 - domains, 47
 - impact of, 34
 - links
 - RSTP (Rapid Spanning Tree Protocol), 74*
 - STP (Spanning Tree Protocol), 71*
 - vPCs (virtual PortChannels), 245*
 - static pinning, 302
 - VDCs (virtual device contexts), 216-217
- farms, servers, 115**
 - configuration, 166
 - round-robin predictors, 116
- Fast Ethernet, 26**
- fast reroute, 340**
- fault isolation, 187, 331**
- fault tolerance**
 - increasing, 240
 - virtual contexts (ACE), 177-178
- FCF (FCoE Forwarder), 505, 510**
- FCIDs (Fibre Channel Identifiers), 413, 420, 429, 508**
- FCIP (Fibre Channel over IP), 453-464**
 - high-availability, 460-462
 - SAN extension with TE, 462
- FC-LS (Fibre Channel Link Services), 422, 476**
- FCoE (Fibre Channel over Ethernet), 186, 493, 504-509**
 - blade servers, 540-542
 - elements, 505-506
 - multihop configuration, 523-535
 - NPIV (N Port ID Virtualization), 532-535
 - SAN extensions, 545
 - unified access server configuration, 515
 - vPCs (virtual PortChannels), 538-540
- FCoE Initialization Protocol.** *See* FIP
- FCoE LEP (FCoE Link End-Point), 506**
- FCoE Link End-Point.** *See* FCoE_LEP
- FEC (Forwarding Equivalence Class), 333**
- Federal Information Processing Standards (FIPS), 227**
- FEX (Fabric Extender), 287, 512**
 - See* FEX
 - connections, 296-299
 - interfaces, 299-301
 - mixed access data centers, 315-317
 - options, 295-296
 - overview of, 291-303
 - redundancy, 301-303

- server access models, 287-291
- status verification, 298, 313
- topologies, 291, 305-315
 - dual-homed*, 309-315
 - straight-through*, 305-308
- fex associate command**, 297
- FHRPs (first-hop redundancy protocols)**, 380-382, 460, 775
 - vPCs (virtual PortChannels), 259-265
- FIB (Forwarding Information Base)**, 87, 92, 190, 334
- Fiber Connectivity (FICON)**, 396
- fiber optics**, 327
 - See also* optical connections
- Fibre Channel**, 396, 410
 - addresses, 413-415
 - classes of service, 420
 - defining, 410-420
 - fabric processes, 420-429
 - FCoE (Fibre Channel over Ethernet), 505
 - flow control, 417-419
 - frames, 415-417
 - layers, 411-412
 - logins, 427-428
 - RSCN (Registered State Change Notification), 426
 - services, 421
 - switches, 506
 - topologies, 412-413
 - zones, 429
- Fibre Channel Adapter Policy option**, 597
- Fibre Channel Identifiers**. *See* FCIDs
- Fibre Channel Link Services (FC-LS)**, 422
- Fibre Channel over Ethernet**. *See* FCoE
- Fibre Channel over IP**. *See* FCIP
- FICON (Fiber Connectivity)**, 396
- fields, VNTag**, 293
- FIF (Fabric Interface)**, 292
- files**
 - answer, 585
 - OVF (Open Virtualization Format), 743
 - storage, 391, 397-398
 - See also* storage
- file systems**, 406-407
 - See also* NFS
- File Transfer Protocol**. *See* FTP
- filters**
 - email, 129
 - reachability, 381
 - VSG (Virtual Security Gateway), 742-753
- Finite State Machine**. *See* FSM
- FIP (FCoE Initialization Protocol)**, 505, 507-509
- FIPS (Federal Information Processing Standards)**, 227
- firewalls**, 111, 737
 - Cisco ASA 1000v, 754-762
 - load balancers, 127-128
 - VSG (Virtual Security Gateway) configuration, 746
- Firewall Service Module**. *See* FWSM
- FireWire**, 396
- firmware policies, UCS (Unified Computing System)**, 633-637
- first-hop redundancy protocols**. *See* FHRPs
- flash memory**, 390
- Flex-Attach, configuration**, 487
- flexibility**, 25, 281
- FlexPod**, 792
- FLOGI (Fabric Login)**, 427, 437, 476, 506
 - frames, 421
 - verification, 513

flooding, 325

MAC (Media Access Control)
addresses, 340

flow control

Fibre Channel, 417-419
PFC (priority-based flow control),
498-499

formatting

See also configuration

FCIDs (Fibre Channel Identifiers),
414

frames

*FCoE (Fibre Channel over
Ethernet), 504*

*VXLANs (Virtual eXtensible
Local Area Networks), 698*

OVF (Open Virtualization Format),
743

security policies, 745-747

VFC (Virtual Fibre Channel) inter-
faces, 525

VSANs (virtual storage-area net-
works), 432-434

Format Unit command, 393**forwarding****frames**

Ethernet, 267

*OTV (Overlay Transport
Virtualization), 354*

vPCs (virtual PortChannels), 243

VRF (Virtual Routing and
Forwarding), 45, 89-91, 199

See also VRF

**Forwarding Equivalence Class. *See*
FEC**

**Forwarding Information Base. *See*
FIB**

Forwarding state (STP), 70

Forwarding Tag (Ftag), 269

**FPMA (Fabric-Provided MAC
Address), 508**

frames

encapsulation, 739

Ethernet, 48, 61, 267

FCoE (Fibre Channel over Ethernet),
504

Fibre Channel, 415-417

OTV (Overlay Transport
Virtualization), 354

parsing, 421

PAUSE, 499

PFC (priority-based flow control),
498

Q-in-Q, 378

SDP (Satellite Discovery Protocol),
296-297

UNPAUSE, 499

vPCs (virtual PortChannels), 243

VXLANs (Virtual eXtensible Local
Area Networks), 698

FSM (Finite State Machine), 625

**FSPF (Fabric Shortest Path First)
protocol, 424-426, 442-444**

Ftag (Forwarding Tag), 269

**FTP (File Transfer Protocol), 167,
750**

F-Trunks, 480

full-mesh extensions, 329

**FWSM (Firewall Service Module),
179, 185**

G

gatekeepers, broadcasts, 365

**Gateway Load Balancing Protocol.
See GLBP**

gateways

default, 11

NAS (Network Attached Storage),
397

VSG (Virtual Security Gateway),
742-753

generic routing encapsulation. *See* GRE

geoclusters, 322-323, 406
 broadcast domain extensions, 325

GET operations, 116-117

Gigabit Ethernet, 26
See also Ethernet
 characteristics, 32
 FCIP (Fibre Channel over IP), 456

GLBP (Gateway Load Balancing Protocol), 12, 380

Globally Unique Identifier. *See* GUID

global resources, 225

Global Server Load Balance. *See* GSLB

governance, certifications, 583

GRE (generic routing encapsulation), 352-354, 803

greenfield data centers, 382-384

groups
 devices, 47
 hierarchies, 14
 ports, 195, 659

growth, attributes of, 25

GSLB (Global Server Load Balance), 775, 777-778

GUID (Globally Unique Identifier), 594

H

HA (high-availability)

FCIP (Fibre Channel over IP), 460-462
 policies, 216, 767

hard disk drives. *See* HDDs

hardware, 19

load balancers, 113
 verification, 190
 x86 servers, 562-572

hashing predictors, 116

HBA (host bus adapters), 393, 509, 585

assigning, 595

vHBAs (virtual HBAs). *See* vHBAs

HDA (Horizontal Distribution Area), 40

HDDs (hard disk drives), 388-389

headers (FabricPath), 268

heartbeat communication, 324-325

hello messages, 38

Hello protocol, 424

hierarchies, groups, 14

HIF (Host Interface), 292, 591

high-availability. *See* HA

High Performance mode, 723-730

history of distributed data centers, 321

1970s-1980s, 321-322

1990s-mid-2000s, 322-324

mid-2000s to present, 324

horizontal cabling, 289

Horizontal Distribution Area (HDA), 40

host bus adapters. *See* HBAs

Host Interface (HIF), 292, 591

hosts, 18

connections, 35

DHCP (Dynamic Host Configuration Protocol), 51

partitioning, 47

software packages, 633

VSM (Virtual Supervisor Module), 662

Hot Standby Router Protocol. *See* HSRP

hot-standby sites, 322, 324

HSRP (Hot Standby Router Protocol), 11-12, 17, 380, 460, 762

HTTP (Hypertext Transfer Protocol),

116, 167
 compression, 130, 134
HTTPS (Secure HTTP), 130
 hub-and-spoke topologies, 329
 hubs
 behavior, 48
 Ethernets, 47
 hybrid clouds, 796
 Hypertext Transfer Protocol. *See*
 HTTP

I

IaaS (Infrastructure as a Service), 795
IBM mainframe virtualization, 10
See also mainframes
ICMP (Internet Control Message Protocol), 116, 163
IDC (International Data Corporation), 2
IDE (Integrated Drive Electronics), 397
 identifiers, VLANs (virtual local-area networks), 377-379, 659
 IDs, bridges, 63
IEC (International Electrotechnical Commission), 29
IEEE (Institute of Electrical and Electronics Engineers), 26
 802.1D, 62
 802.1Q, 54
 XXXX802.1Qaz. *See* ETS
 XXXX802.1Qbb. *See* PFC
 802.1w, 72
 1394 (FireWire), 396
IETF (Internet Engineering Task Force), 12, 130
 imaging, 585
 implementation
 cloud computing, 797-798
 partitioning, 187-188
 VM-FEX (Virtual Machine Fabric Extender), 706
importing plug-ins, 711
INCITS (International Committee for Information Technology Standards), 392
 industrializing server provisioning, 637-653
 inflection point positions, 119
 information security, 735
infrastructure
 domains, 585, 594
 IVR (Inter-VSAN Routing), 465-467
 layers, 189
 UCS (Unified Computing System), 586
 VM-FEX (Virtual Machine Fabric Extender), 705
Infrastructure as a Service. *See* IaaS
Ingress Tunnel Router (ITR), 779
 inheritance, port profiles, 665
 initializing fabric, 421-424
 initiating SSLs (Secure Sockets Layers), 132
 input/output. *See* I/O
Inquiry command, 393
installation
 access switches, 289
 Cisco ASA 1000v, 755-757
 operating systems, 620-623
 VEM (Virtual Ethernet Module), 671
 VSG (Virtual Security Gateway), 743-745
 VSM (Virtual Supervisor Module), 671
instances (STP), 74-78
Institute of Electrical and Electronics Engineers. *See* IEEE
Integrated Drive Electronics (IDE), 397

integration

- networks, 660
- UCS VMware Integration wizard, 711-716
- virtual contexts (ACE), 156-161

Intel, 26**Intelligent Platform Management Interface. *See* IPMI****interconnections, 18****Interface Manager, 189****interfaces**

- allocation, 190
- API (application programming interface), 661
- ATAPI (ATA Packet Interface), 397
- attachments, 343
- backup, 495
- BVIs (Bridge-group Virtual Interfaces), 145, 158
- CLIs (command-line interfaces), 20, 142
- clusters, 700
- DCI (Data Center Interconnect), 329
- dynamic, provisioning, 664-666
- EMIs (electromagnetic interfaces), 28
- Ethernet, 591
 - 10/10 configuration, 197-198*
 - LACP modes, 237*
- Fabric Extender, 299-301
- FabricPath, 273
- FCoE (Fibre Channel over Ethernet), 186
- Fibre Channel, 413
- FIF (Fabric Interface), 292
- Gigabit Ethernet, 456
- HIF (Host Interface), 292
- internal, 357
- IPMI (Intelligent Platform Management Interface), 609

join, 357

- Layer 2 connections, 343
- LIF (Logical Interface), 292
- loops, verifying status of, 192
- maps, 771
- migration, 687
- MPC (Modular Policy CLI), 162
- MTUs (maximum transmission units), 338
- multiple, 54
- NPV (N_Port Virtualization), 480
- out-of-band management switches, 143
- overlay, 358, 369
- PFC (priority-based flow control), 498
- policy maps, 162
- private VLANs, 79
- production switch configuration, 142
- promiscuous access, 81
- public/private, 495
- Q-in-Q configuration, 378
- redundancy, 37, 301-303
- SCSI (Small Computer Systems Interface), 392-396, 410
- servers, 591
- SVIs (Switch Virtual Interfaces), 58, 145, 669, 792
- switches
 - port IDs, 65*
 - RSTP (Rapid Spanning Tree Protocol), 72*
- trunk configuration, 55
- uplinks, 591
- vCons (virtual network interface connections), 602
- VFC (Virtual Fibre Channel), 512, 525
- VIF (Virtual Interface), 292

- virtual
 - placement*, 602-603
 - verification*, 690
- VLANs (virtual local-area networks), 52
- VNMC (Virtual Network Manager Center), 743
- VSG (Virtual Security Gateway), 744
- VSM (Virtual Supervisor Module), 662, 682
- interference, troubleshooting, 28
- Intermediate System-to-Intermediate System. *See* IS-IS
- internal interfaces, 357
- internal routing (Layer 2 extensions), 380-382
- internal security zones, 225
- Internal Spanning Tree. *See* IST
- International Committee for Information Technology Standards (INCITS), 392
- International Data Corporation (IDC), 2
- International Electrotechnical Commission (IEC), 29
- International Organization for Standardization (ISO), 29
- Internet boom (1990s), 5
- Internet Control Message Protocol. *See* ICMP
- Internet Engineering Task Force. *See* IETF
- Internet Protocol. *See* IP
- Internet SCSI. *See* iSCSI
- Internet service providers. *See* ISPs
- interoperability between MST/STP switches, 78
- interswitch connections, 234
- Inter-Switch-Link. *See* ISL
- Inter-VSAN Routing. *See* IVR
- intra-rack cabling, 40
- intrusion prevention services. *See* IPS
- I/O (input/output), 388, 493
 - channel architecture, 396
 - convergence, 495-497
 - HDDs (hard disk drives), 388
- IP (Internet Protocol), 26, 46
 - addresses
 - assigning*, 59
 - Layer 4/Layer 7 switches*, 120
 - management*, 610
 - virtualization clusters*, 668
 - VSG (Virtual Security Gateway)*, 752
 - connections, testing, 200
 - FCIP (Fibre Channel over IP), 453-464
 - Layer 2 extensions, 352-377
 - readdressing, 325
 - routing, 87
 - subnets, VLAN association to, 56-58
- IPMI (Intelligent Platform Management Interface), 609
- IPS (intrusion prevention services), 111, 225
- IPSec (IP Security), 774
- iSCSI (Internet SCSI), 396, 495
 - vNICs (virtual NICs), 604
- IS-IS (Intermediate System-to-Intermediate System), 87
 - FabricPath, 330
- ISL (Inter-Switch-Link), 56
 - isolation, 424
 - NPV (N_Port Virtualization), comparing to, 475
- islands, SANs (storage area networks), 430-432
- ISO (International Organization for Standardization), 29
- isolation
 - fault, 331

ISL (Inter-Switch-Link), 424
 Layer 2 extensions, 326
 private VLANs, 78-83
 traffic, 47
 ISPs (Internet service providers), 35
 IST (Internal Spanning Tree), 77
 ITR (Ingress Tunnel Router), 779
 IVR (Inter-VSAN Routing), 453, 464-473
 infrastructure, 465-467
 transit VSANs (virtual storage-area networks), 472
 zones, 467-472

J

Java DataBase Connectivity. *See* JDBC
 JBODs (Just a Bunch of Disks), 389
 JDBC (Java DataBase Connectivity), 398
 join interfaces, 357
 Just a Bunch of Disks. *See* JBODs

K

keepalive links
 connections, 248
 peers, 245
 Kerberos, 398
 kernels, 189
 KVM tool, 620

L

Label Distribution Protocol. *See* LDP
 label edge routers. *See* LERs
 Label Forwarding Information Base. *See* LFIB
 labels

binding, 337, 353
 MPLS (Multiprotocol Label Switching), 47, 332
 label switch routers. *See* LSRs
 LACP (Link Aggregation Control Protocol), 237, 686
 LAN Emulation (LANE), 803
 LANs (local-area networks), 11, 49
 SAN management separation, 546-555
 VXLANs (Virtual eXtensible Local Area Networks), 697-705
 LAN tab (UCS Manager), 591
 Layer 2
 distributed data centers. *See* distributed data centers
 extensions, 319
FabricPath, 330-332
internal routing, 380-382
IP (Internet Protocol), 352-377
MPLS (Multiprotocol Label Switching), 332-351
OTV (Overlay Transport Virtualization), 354-377
overview of, 324-327
solutions, 382-384
VLAN (virtual local-area network) identifiers, 377-379
vPCs (virtual PortChannels), 328-330
VPLS (Virtual Private LAN Service), 342-351
 interface connections, 343
 multipathing
with FabricPath, 43
vPCs (virtual PortChannels), 265-280
 VDCs (virtual device contexts), 190
 Layer 3
 PortChannel configuration, 236

- routing tables, 88
- VDCs (virtual device contexts), 190
- VLANs (virtual local-area networks), 58-61
- Layer 4**
 - connection management, 122
 - load-balancing policy maps, 168
 - switches, 120-121
- Layer 7**
 - class maps, 167
 - connection management, 122
 - load-balancing policy maps, 169
 - switches, 113, 120-121
- layers**
 - access, 35, 542-545
 - aggregation, 36
 - distribution, 35
 - FCoE (Fibre Channel over Ethernet), 505
 - Fibre Channel, 411-412
 - infrastructure, 189
 - networks, 35
 - parameters, 112
 - physical, 27
 - coaxial cables*, 27
 - optical fiber*, 29
 - standards*, 29
 - server access (Unified Fabric designs), 536-542
 - SSL (Secure Sockets Layer), 130-133
 - TLS (Transport Layer Security), 130
- LBAs (Logical Block Addresses), 394**
- LDP (Label Distribution Protocol), 334**
- Learning state (STP), 70**
- least-connections predictors, 116**
- least-loaded predictors, 117**
- legacy extension disabling, 371**
- length, transfer, 394**
- LERs (label edge routers), 333**
- LFIB (Label Forwarding Information Base), 334**
- libraries**
 - tape, 390-391
 - VTL (virtual tape library), 403
- LID (Local ID), 268**
- LIF (Logical Interface), 292**
- light propagation, 31**
- limitations**
 - of connections, 118
 - of data centers, 5
 - of monitor sessions (VDCs), 204
 - of scalability, 17
 - weights, 118
- linecards**
 - sharing, 191
 - VDCs (virtual device contexts), 195
- Link Aggregation Control Protocol. See LACP**
- Link Layer Discovery Protocol. See LLDP**
- links**
 - aggregation, 234-240
 - failures
 - RSTP (Rapid Spanning Tree Protocol)*, 74
 - STP (Spanning Tree Protocol)*, 71
 - vPCs (virtual PortChannels)*, 245
 - STP (Spanning Tree Protocol), 232-234
 - Virtual Link, 506
 - virtual networking, 684-688
 - VSL (virtual switch link), 241
- link-state records. See LSRs**
- link-state updates. See LSUs**
- Linnaeus, Carl, 21**
- Linux, 9**

LISP (Location ID/Separation Protocol), 775, 779-781, 803

Listening state (STP), 70

lists, access, 164

See also ACLs

LLDP (Link Layer Discovery Protocol), 502

load balancers

address translation, 124-127

applications, 127-130

applying, 111-134

concepts, 115-120

firewalls, 127-128

offloading servers, 130-134

performance, 135-136

policies

configuration, 132

maps, 168

proliferation in data centers, 135-139

resources, 693

reverse proxies, 128-130

scalability, 135

security policies, 136-137

SSL offloads, 132

suboptimal traffic, 137

transparent mode, 126-127

virtual contexts (ACE), 164

load balancing, 43, 58, 225

local-area networks. *See* LANs

Local Disk option, 604

Local ID. *See* LID

localization, traffic, 382

Location ID/Separation Protocol. *See* LISP

Logical Block Addresses. *See* LBAs

Logical Interface (LIF), 292

logical partitioning, 42-43

logical topologies, 213

logical unit numbers. *See* LUNs

logical units. *See* LUs

Logical Volume Manager. *See* LVM
login

FIP (FCoE Initialization Protocol), 507

FLOGI (Fabric Login). *See* FLOGI
user configuration, 176

logins

Fibre Channel, 427-428

VLAN commands, testing, 86

Login Server service, 421

Long Distance vMotion, 693

longevity of Ethernets, 26

loop-free U topologies, 38

loops

arbitrated, 412

avoidance, 326, 330

broadcasts (Ethernets), 62

Cisco Nexus 1000v switches, 685

detection, 193

interfaces, verifying status of, 192

OTV (Overlay Transport
Virtualization), 365-373

STP (Spanning Tree Protocol), 37, 70

loss, frames, 498

LSRs (label switch routers), 333

LSRs (link-state records), 424

LSUs (link-state updates), 424

LUNs (logical unit numbers), 393,
399

masking, 585

SCSI storage volumes, 439

virtualization, 404-406

LUs (logical units), 393

LVM (Logical Volume Manager), 392

M

MAC (Media Access Control), 51

addresses, 581

*BPDU frames, 65**FabricPath, 331**flooding, 340**FPMA (Fabric-Provided MAC Address), 508**learned in MPLS, 346**OTV (Overlay Transport Virtualization), 355**service profiles, 587**sharing resources, 85**SPMA (Server-Provided MAC Address), 508**tables, 691*

Ethernet, 505

pinning, 686

security, 601

magnetic tapes, 391**Main Distribution Area (MDA), 40****mainframes, 560-561**

storage access, 396-397

virtualization, 10

zSeries, 397

main memory, 8**maintenance policies, 606****Manage Hosts vCenter wizard, 606****management**

ASDM (Adaptive Security Device Manager), 754

CIMC (Cisco Integrated Management Controller), 588

Cisco ASA 1000v, 755

Cisco Nexus 1000v architecture, 661

CMC (chassis management circuit), 588

connections, 122-124

consolidation, 44

EEM (Event Embedded Manager),

225

Fabric Extenders, 291

firmware packages, 633

interfaces, 495

IP (Internet Protocol) addresses, 610

IPMI (Intelligent Platform Management Interface), 609

LUNs (logical unit numbers), 404

planes, 20

SNMP (Simple Network Management Protocol), 20, 85

traffic

*allowing to virtual contexts (ACE), 162-163**NPV (N_Port Virtualization), 482-486*

UCS (Unified Computing System), 578, 588

VDCs (virtual device contexts), 185, 214-224

*operations, 214-216**out-of-band, 217-222**process failures, 216-217**RBAC (Role-Based Access Control), 222-224*

virtual contexts (ACE), 163,171-176

VRF (Virtual Routing and Forwarding), 98-101

VSG (Virtual Security Gateway), 751

VSM (Network Services Manager), 792

Management Server service, 422**Manage Plug-ins option, 711****Map-Resolver (MR), 779****maps**

classes

*Layer 7, 167**VIP (virtual IP), 169*

interfaces, 771

policies, 162

- load balancers*, 168
- multimatch*, 169
- QoS (quality of service), 225
- vCons (virtual network interface connections), 602
- Map-Server (MS)**, 779
- masking LUNs (logical unit numbers)**, 585
- Massive Parallel Processing**. *See* **MPP**
- maximum transmission units**. *See* **MTUs**
- MBps (megabytes per second)**, 568
- MC-LAG (Multi-Chassis Link Aggregation Group)**, 351
- MDA (Main Distribution Area)**, 40
- MDS (Multilayer Data Switch)**, 410, 424, 456
- MDS-CORE routing tables**, 442
- measurements**
 - inflection points, 119
 - services, 794
- media**
 - Ethernet protocol, 27
 - Gigabit Ethernet, 32
- Media Access Control**. *See* **MAC**
- megabytes per second**. *See* **MBps**
- memory**
 - allocation, 210
 - caches, 9
 - estimation, 211
 - flash, 390
 - modules, 568
 - NUMA (Non-Uniform Memory Access), 566
 - ranges, 210
 - VDCs (virtual device contexts), 203
 - virtual, 8-9
 - x86
 - architecture*, 562
 - servers*, 566-569
- merging data centers**, 731-733
- Message and Translation System (MTS)**, 662
- messages**
 - ACCEPT, 427
 - CNMs (congestion notification messages), 503
 - hello, 38
 - SMB (Server Message Block), 398
- metadata**, 391
- Metcalfe, Robert**, 26
- Microsoft Windows**, 9
- migration**
 - Cisco Nexus 1000v switches, 693-697
 - connections, 289
 - OTV (Overlay Transport Virtualization), 366
 - physical interfaces, 687
 - servers, 325
 - UCS Manager, 625
 - virtual LUNs, 406
 - VM-FEX (Virtual Machine Fabric Extender), 720-722
 - VM-Nomad, 770
 - VMs (virtual machines) to VM-FEX (Virtual Machine Fabric Extender), 716-720
 - VSM (Virtual Supervisor Module) interfaces, 682
- mirroring RAM (random-access memory)**, 569
- misconceptions about VLANs (virtual local-area networks)**, 56-61
- mismatches, native VLANs**, 84
- mixed access data centers (Fabric Extender)**, 315-317
- MMF (multimode fiber)**, 29, 396
- mobile devices**, 168
 - See also* **devices**
- models**

- converged access, 542
- converged aggregation, 543
- OSI (Open Systems Interconnection), 411
- point-to-point server access, 537
- server access, 287-291
- VSG (Virtual Security Gateway) management, 751
- modes**
 - ON, 238
 - address translation, 124
 - communication (Cisco Nexus 1000v switches), 663-664
 - End-Host, 592
 - High Performance, 723-730
 - transparent, 126-127
- modification**
 - See also* configuration
 - resource allocation, 156
 - service profile associations, 614
- modularization, software, 5**
- Modular Policy CLI. *See* MPC**
- modules**
 - hardware, verification, 190
 - memory, 568
 - VEM (Virtual Ethernet Module), 661, 739-740
 - VSM (Virtual Supervisor Module), 661
- monitoring**
 - performance, 111
 - policies, 611
 - sessions, 203
- Moore's Law, 564**
- motherboards (x86 architecture), 564**
- MPC (Modular Policy CLI), 162**
- MPLS (Multiprotocol Label Switching), 47**
 - connections, 353
 - EoMPLS (Ethernet over MPLS), 338-342
 - Layer 2 extensions, 332-351
 - over GRE (generic routing encapsulation), 352-354
 - overview of, 333-338
- mpls ip command, 335**
- MPP (Massive Parallel Processing), 566**
- MR (Map-Resolver), 779**
- MS (Map-Server), 779**
- MSFC (Multilayer Switching Feature Card), 145**
- MST (Multiple Spanning Tree), 76-77, 187**
 - configuration, 201
 - verification, 202
- MTS (Message and Translation System), 662**
- MTU option, 601, 803**
- MTUs (maximum transmission units), 338, 460, 504, 740**
- Multi-Chassis Link Aggregation Group (MC-LAG), 351**
- multihoming, 326, 328, 365-373**
- multihop FCoE (Fibre Channel over Ethernet) configuration, 523-535**
- Multilayer Data Switch. *See* MDS**
- Multilayer Switching Feature Card. *See* MSFC**
- multimatch policy maps, 169**
- multimode fiber (MMF), 29, 396**
- multipathing**
 - with FabricPath, 43
 - vPCs (virtual PortChannels), 265-280
- multiple connections, STP (Spanning Tree Protocol), 233**
- multiple data centers, 319**
 - See also* data centers
- multiple interfaces, 54**
- Multiple Spanning Tree. *See* MST**

multiple VLANs, 59

Multiprotocol Label Switching. *See* MPLS

multisite interconnections (FabricPath), 331

multitenant data centers, virtual contexts (ACE), 179-181

N

N7K switches, 94, 190, 519

connections, 524

private VLANs, 82

NAA (network address authority), 414

names

VDCs (virtual device contexts), 198

VMs (virtual machines), 752

WWNs (World Wide Names), 581

NAS (Network Attached Storage), 397

NAT (Network Address Translation), 124, 466, 762

Dual NAT, 125-126

Server NAT, 124-125

National Institute of Standards and Technology (NIST), 794

native VLANs, 84

neighborhood information (OSPF), 95

nesting, 199-202

NetApp, 403

NetBIOS, 398

network address authority. *See* NAA

Network Address Translation. *See* NAT

Network Attached Storage. *See* NAS

Network Control Policy option, 601

Network File System (NFS), 398

network interface cards. *See* NICs

networks

adapters, convergence, 509

availability, 37

cloud computing, 800

connections (ACE 4710), 141

convergence, 495-497

design, 289-291

See also design

Ethernets. *See* Ethernets

evolution, 25

extensions, 44

integration, 660

LANs (local-area networks), 11, 49

load balancers, 111-134

logical partitioning, 42-43

MPLS (Multiprotocol Label Switching), 335

ONE (Cisco Open Network Environment), 804

OpenStack, 801-802

overlapping, 87

overlays, 802-804

partitioning, 47

PODs (pool of devices), 281-284

SDNs (Software-Defined Networks), 800-801

segmentation, 103-104

services, 735

applications, 111

data paths, 738-740

definitions, 735-738

profiles, 599-608

vPath-enabled, 740-771

VSG (Virtual Security Gateway), 742-753

simplification, 43

SNMP (Simple Network Management Protocol), 20, 85

topologies, 34-41

layers, 35-

- physical layouts*, 39-40
 - virtualization, 8, 42-44
 - virtual networking, 658-661
 - NX-OS features*, 688-693
 - VM-FEX (Virtual Machine Fabric Extender)*, 705-707, 720-722
 - VXLANs (Virtual eXtensible Local Area Networks)*, 697-705
 - VLANs (virtual local-area networks), 45
 - VPNs (Virtual Private Networks), 47
 - WANs (wide-area networks), 4
 - Network Services Manager**. *See* NSM
 - new cps (connections per second)**, 135
 - Nexus 1000v switches**
 - See also* switches
 - application acceleration, 763-771
 - connections, 741
 - frame encapsulation, 739
 - virtual networking topology, 741
 - vPath configuration, 748, 758
 - Nexus 7000 switches**, 92
 - See also* switches
 - commands, 61
 - OTV (Overlay Transport Virtualization), 354-377
 - port groups, 195
 - NFS (Network File System)**, 398, 495, 763
 - NICs (network interface cards)**, 54, 509
 - physical NIC teaming, 660
 - teaming, 238
 - virtual NICs. *See* vNICs
 - NIST (National Institute of Standards and Technology)**, 794
 - nodes**, 322
 - See also* servers
 - status, 751
 - virtual service, 740
 - nondisruptive reconfiguration of fabric**, 423
 - Non-Uniform Memory Access**. *See* NUMA
 - no resource allocation**, 17
 - notification**
 - CNMs (congestion notification messages), 503
 - QCN (Quantized Congestion Notification), 503
 - NPIV (N_Port ID Virtualization)**, 476
 - configuration, 476-482
 - FCoE (Fibre Channel over Ethernet), 532-535
 - N_Port ID Virtualization (NPIV)**, 476
 - N_Port Virtualization**. *See* NPV
 - NPV (N_Port Virtualization)**, 453, 473-490
 - blade servers, 488
 - port WWN virtualization deployment, 486-488
 - traffic management, 482-486
 - NT LAN Manager**. *See* NTLM
 - NTLM (NT LAN Manager)**, 398
 - NUMA (Non-Uniform Memory Access)**, 566
 - NX-OS switches**, 63, 188
 - virtual networking, 688-693
 - VSANs (virtual storage-area networks), 432
-
- ## O

- ODBC (Open DataBase Connectivity)**, 398
 - offloading servers, load balancers**, 130-134
 - on-demand self-service**, 794

ONE (Cisco Open Network Environment), 804

one-armed design, virtual contexts (ACE), 160-161

ONF (Open Networking Foundation), 800

online migrations

Cisco Nexus 1000v switches, 693-697

VM-FEX (Virtual Machine Fabric Extender), 720-722

ON mode, 238

ONS 1500, 328

See also DWDM

OOO (out-of-order), 148, 268

Open DataBase Connectivity. *See* ODBC

OpenFlow, 801

Open Networking Foundation (ONF), 800

Open Shortest Path First. *See* OSPF

Open Software Foundation (OSF), 594

OpenStack, 801-802

Open Systems Interconnection. *See* OSI

Open Virtualization Appliance. *See* OVA

Open Virtualization Format. *See* OVF
operating systems, 19

converged network adapters, 509

installation, 620-623

NX-OS. *See* NX-OS switches

server domains, 585

sharing, 10

operational areas, 5-7

operational policies (service profiles), 608-612

operations, VDC (virtual device context) management, 214-216

optical connections, Ethernet extensions over, 327-332

optical fiber, 29-32, 34

optimization

cables, 44

data centers, 231

responses, 130

STP (Spanning Tree Protocol), 72-73

vWAAS (Virtual Wide Area Application Services),

options

Any Configuration, 595

Edit Virtual Machine Settings, 690

Ethernet Adapter Policy, 601

Fabric Extenders, 295-296

Fibre Channel Adapter Policy, 597

Local Disk, 604

Manage Plug-ins, 711

MTU, 601

Network Control Policy, 601

Persistent Binding, 597

Pin Group, 597, 601

Protection Configuration, 596

QoS Policy, 597, 601

vHBAs (virtual HBAs), 604

Virtual Media, 604

vNICs (virtual NICs), 604

Organizational Unique Identifier. *See* OUI

origins of virtualization, 8-13

OSF (Open Software Foundation), 594

OSI (Open Systems Interconnection), 411

OSPF (Open Shortest Path First), 87, 94, 95, 334

OTV (Overlay Transport Virtualization), 319, 352, 354-377, 359-364, 803

OUI (Organizational Unique Identifier), 414

out-of-band management

interface configuration, 143
 VDCs (virtual device contexts), 217-222

out-of-order (OOO), 148, 268

OVA (Open Virtualization Appliance), 671

overlapping

addresses, 87-89
 FPMAs (Fabric-Provided MAC Addresses), 509

overlays

interfaces, 358, 369
 networks, 802-804

Overlay Transport Virtualization. *See* OTV

oversubscription ratios, 37

OVF (Open Virtualization Format), 743, 755

P

PaaS (Platform as a Service), 795

packets

ATAPI (ATA Packet Interface), 397
 MPLS (Multiprotocol Label Switching), 332

Palo Alto Research Center (PARC), 26

parallel expansion busses, 570

parameters

ASA 1000v OVF, 755
 layers, 112
 VSM (Virtual Supervisor Module), installation, 671

PARC (Palo Alto Research Center), 26

parent switches

See also switches
 connections, 296-299
 Fabric Extenders, 310

parity, 569

parsing frames, 421

partitioning

implementation, 187-188
 networks, 42-43, 47
 VDCs (virtual device contexts), 185

PAT (Port Address Translation), 124

paths. *See* vPath

behavior, 47
 default path cost for switches, 64
 network services, 738-740

patterns, traffic, 245

PAUSE frames, 499

Payment Card Industry (PCI), 227

PBR (policy-based routing), 127, 765

PCI (Payment Card Industry), 227

PCP (Priority Code Point), 53

Pearlman, Radia, 62

PE (provider edge) devices, 338

peers, keepalive links, 245, 248

Peer-to-Peer Remote Copy (PPRC), 403

Pelissier, Joe, 509

Penultimate Hop Popping. *See* PHP performance

High Performance mode, 723-730
 load balancers, 111, 135-136
 monitoring, 111

peripherals, 4

permissions, access, 391

Persistent Binding option, 597

Persistent Lempel-Ziv (PLZ) compression, 764

Per VLAN Spanning Tree. *See* PVST

Per VLAN Spanning Tree Plus. *See* PVST+

PFC (priority-based flow control), 498-499

PHP (Penultimate Hop Popping), 334

physical connections, 495

ACE (Application Control Engine), 141-145

physical formats (x86 servers), 571-572

physical interfaces, migration, 687

physical layers, 27

coaxial cables, 27

optical fiber, 29

standards, 29

physical layouts, networks, 39-40

physical NIC teaming, 660

physical servers

See also servers

activation, 581

migration (UCS Manager), 625

provisioning, 581

physical-to-logical storage, 8

Pick-up Truck Access Method (PTAM), 322

PING, 166

Pin Group option, 597, 601

pinning

MAC (Media Access Control), 686

static, 301

placement of virtual interfaces, 602-603

planes, 20, 98-101

plans, disaster recovery, 321

Platform as a Service. *See* PaaS

PLOGI (Port Login), 427

plug-ins

exporting, 712

importing, 711

XML (Extensible Markup Language), 742

PLZ (Persistent Lempel-Ziv) compression, 764

PODs (pool of devices), 281-284, 328, 790

point-to-point

ports, 72, 412

server access models, 537

policies

BIOS (basic input/output system) configuration, 627-633

firmware (UCS), 633-637

HA (high-availability), 216

load balancers

configuration, 132

maps, 168

maintenance, 606

maps, 162

MPC (Modular Policy CLI), 162

multimatch, 169

operational (service profiles), 608-612

security, 737

Cisco ASA 1000V, 761

formatting, 745-747

load balancers, 136-137

services, 170

UCS (Unified Computing System), 626-633

verification, 752

vPath, 739

policy-based routing. *See* PBR

pooling, 16, 405, 639-640

cloud computing, 786

resources, 794

servers, 649-653

virtualization, 116

pool of devices. *See* PODs

Port Address Translation. *See* PAT

PortChannels, 141, 203

Cisco Nexus 1000v switches, 687

configuration, 303, 306

cross-switch, 240-241

FCIP (Fibre Channel over IP), 462

FSPF (Fabric Shortest Path First) protocol, 426

Layer 3 configuration, 236

- load-balancing methods, 225
- MDS-B, configuration on, 436
- STP (Spanning Tree Protocol) configuration, 236
- straight-through topologies, 306
- VFC (Virtual Fibre Channel) configuration, 528-531
- vpc-hm (Virtual PortChannel Host-Mode), 686
- vPCs (virtual PortChannels), 231, 241-265
 - See also* vPCs (virtual PortChannels)
- Port Login (PLOGI), 427**
- ports**
 - access, 390
 - connections (VDCs), 192
 - Designated Ports Selection phase, 68
 - Ethernets, 26, 72
 - groups, 659
 - IDs, 65
 - Nexus 7000 switches, 195
 - NPV (N Port Virtualization), 453, 473-490
 - point-to-point, 412
 - private VLANs, 79
 - profiles
 - Cisco Nexus 1000v switches, 664-666*
 - vMotion, 682*
 - VSG data interfaces, 744*
 - promiscuous ports, 81
 - redirection, 126
 - roles, 73
 - Root Port Selection phase, 68
 - Smart Port Macros, 664
 - states, 70-72
 - switch connections, 69
 - translation, 124
 - types (Fibre Channel), 412-413
 - WWNs (World Wide Names), 486-488
- port world wide name (pWWN) addresses, 465**
- power**
 - control policies, 611
 - distribution, 7
 - supply redundancy, 390
- PPRC (Peer-to-Peer Remote Copy), 403**
- predictability, 25, 281**
- predictors, 116, 167**
- preprovisioning (Fabric Extender), 311**
- primary storage, 8, 388**
- printers, 4**
- prioritization, 37**
- priority-based flow control. *See* PFC**
- Priority Code Point. *See* PCP**
- private clouds, 796**
- private interfaces, 495**
- private VLANs, 78-83**
- PRLI (Process Login), 427**
- probes, 116, 136, 164-167**
- processes**
 - failures (VDCs), 216-217
 - sharing, 189
- Process Login (PRLI), 427**
- production switches, 142**
- profiles**
 - FCIP (Fibre Channel over IP), 460
 - ports
 - Cisco Nexus 1000v switches, 664-666*
 - vMotion, 682*
 - VSG data interfaces, 744*
 - services, 581
 - associating to servers, 612-619*
 - boot order, 604-605*

- building*, 588-623
- identifying*, 594
- industrializing server provisioning*, 637-653
- installing operating systems*, 620-623
- maintenance policies*, 606
- network definitions*, 599-608
- operational policies*, 608-612
- policies*, 626-633
- seasonal workloads*, 653-654
- server assignment*, 606-607
- storage definitions*, 595-598
- templates*, 640-649
- UCS (Unified Computing System)*, 586-588
- virtual interface placement*, 602-603
- vNICs (virtual NICs)*, 707-709
- promiscuous access interfaces**, 81
- prompts**
 - CLI (command-line interface), 198
 - EXEC, 147
- propagation, light**, 31
- properties, security**, 659
- proprietary protocols**, 392
- protected space**, 188
- Protection Configuration option**, 596
- protocols**
 - ARP (Address Resolution Protocol), 49, 701
 - BGP (Border Gateway Protocol), 87, 334
 - BOOTP (Bootstrap Protocol), 51
 - CDP (Cisco Discovery Protocol), 601
 - DCBX (Data Center Bridging eXchange), 501-503, 518
 - DHCP (Dynamic Host Configuration Protocol), 51, 585, 762
 - EIGRP (Enhanced Interior Gateway Routing Protocol), 87
 - Ethernet, 25-26
 - coaxial cables*, 27-28
 - direct-attach twinaxial cables*, 32-33
 - media*, 27
 - optical fiber*, 29-32
 - overview of*, 26-34
 - timelines*, 33-34
 - twisted-pair cables*, 28-29
 - FHRPs (first-hop redundancy protocols), 380-382, 460, 775
 - FHRPs (first-hop routing protocols), 259-265
 - FIP (FCoE Initialization Protocol), 505, 507-509
 - FSPF (Fabric Shortest Path First), 424-426
 - FTP (File Transfer Protocol), 167, 750
 - GLBP (Gateway Load Balancing Protocol), 12, 380
 - Hello, 424
 - HSRP (Hot Standby Router Protocol), 11-12, 17, 380, 460, 762
 - HTTP (Hypertext Transfer Protocol), 130, 134, 167
 - ICMP (Internet Control Message Protocol), 116, 163
 - IP (Internet Protocol), 26, 46
 - FCIP (Fibre Channel over IP)*, 453-464
 - Layer 2 extensions*, 352-377
 - readdressing*, 325
 - routing*, 87
 - ISL (Inter-Switch-Link), 56
 - LACP (Link Aggregation Control Protocol), 237, 686

- LDP (Label Distribution Protocol), 334
- LISP (Location ID/Separation Protocol), 775, 779-781, 803
- LLDP (Link Layer Discovery Protocol), 502
- proprietary, 392
- RIP (Routing Information Protocol)
- routing, 92-98
- RSTP (Rapid Spanning Tree Protocol), 72
- RTSP (Real-Time Streaming Protocol), 167
- SDP (Satellite Discovery Protocol), 296
- SIP (Session Initiation Protocol), 167
- SNMP (Simple Network Management Protocol), 20, 85, 117
- SRP (Satellite Registration Protocol), 296, 298
- standard, 392
- STP (SATA Tunneling Protocol), 397
- STP (Spanning Tree Protocol), 37
 - convergence*, 66
 - DCI (Data Center Interconnect)*, 326
 - design*, 232
 - Fabric Extender interfaces*, 299-301
 - FabricPath*, 272-276, 330
 - instances*, 74-78
 - link utilization*, 232-234
 - optimization*, 72-73
 - port states*, 70-72
 - scalability*, 326
 - status verification*, 274
 - switches*, 63
 - VDCs (virtual device contexts)*, 192
 - VLANs (virtual local-area networks)*, 61-78
 - vPCs (virtual PortChannels)*, 254-259
- TCP (Transmission Control Protocol), 120
 - offloading servers*, 130, 133-134
 - port redirection*, 126
 - timeouts*, 123
- TCP/IP (Transmission Control Protocol/Internet Protocol), 11
- TFTP (Trivial File Transfer Protocol), 750
- UDP (User Datagram Protocol), 120
 - port redirection*, 126
 - virtual contexts (ACE)*, 164
- ULP (upper-layer protocol), 412
- VRRP (Virtual Router Redundancy Protocol), 12, 380, 460
- WCCP (Web Cache Control Protocol), 765
- provider edge. See PE**
- provisioning**
 - automation, 581
 - dynamic interfaces, 664-666
 - PODs (pool of devices), 790
 - servers, 581
 - activation*, 583-585
 - industrializing*, 637-653
- pseudowires**, 343
- PTAM (Pick-up Truck Access Method)**, 322
- public clouds**, 796
- public interfaces**, 495
- PVST (Per VLAN Spanning Tree)**, 74
- PVST+ (Per VLAN Spanning Tree Plus)**, 74
- pWWN (port world wide name) addresses**, 465

Q

QCN (Quantized Congestion Notification), 503

Q-in-Q frame encapsulation, 378

QoS (quality of service), 340

- maps, 225
- Policy option, 597, 601
- service profiles, 587

QSFP+ (Enhanced Quad Small Form Factor Pluggable), 32

quality of service. *See* QoS

Quantized Congestion Notification. *See* QCN

R

racks, 657

Rackspace, 794

rack units. *See* RUs

RADIUS (Remote Authentication Dial-In User Service), 167

RAID (Redundant Array of Independent Disks), 399-401, 596

RAM (random-access memory), 390, 566, 569

ranges, memory, 210

rapid elasticity, 794

Rapid PVST+ (Rapid Per VLAN Spanning Tree Plus), 74

Rapid Spanning Tree Protocol. *See* RSTP

RAS (Reliability, Availability, and Serviceability), 608

rate limiter (RL), 503

ratios, oversubscription, 37

RBAC (Role-Based Access Control), 171

- UCS (Unified Computing System), 586
- VDCs (virtual device contexts), 222-224

RCF (Reconfigure Fabric) frame, 423

RDIMMs (Registered DIMMs), 569

RDIs (request domain identifiers), 423

reachability, filtering, 381

Read Capacity command, 393

Read command, 393

readdressing, IP (Internet Protocol), 325

real server configuration, 164-167

Real-Time Streaming Protocol. *See* RTSP

Receive Side Scaling (RSS), 725

Reconfigure Fabric (RCF) frame, 423

records

- LSRs (link-state records), 424
- storage access, 398-399

recovery, disaster recovery plans, 321

Recovery Point Objective (RPO), 321

Recovery Time Objective (RTO), 321

redirection

- ports, 126
- vPath, 739

Reduced Instruction Set Computing. *See* RISC servers

redundancy

- CRC (cyclic redundancy check), 417
- DRE (Data Redundancy Elimination), 764
- Fabric Extender interfaces, 301-303
- FHRPs (first-hop redundancy protocols), 380-382, 460, 775
- interfaces, 37
- power supplies, 390
- storage virtualization, 402
- VPLS (Virtual Private LAN Service), 348
- VRRP (Virtual Router Redundancy Protocol), 460

- VSG (Virtual Security Gateway), 744
- Redundant Array of Independent Disks.** *See* RAID
- reflection, 29
- Registered DIMMs.** *See* RDIMMs
- Registered State Change Notification.** *See* RSCN
- registration**
 - VNMC (Virtual Network Manager Center), 756
 - VSM (Virtual Supervisor Module), 744
- reliability**
 - attributes of, 25
 - of servers, 559
- Reliability, Availability, and Serviceability (RAS), 608**
- reload vdc command, 215
- remapping addresses, 17
- Remote Authentication Dial-In User Service.** *See* RADIUS
- Remote Shell.** *See* RSH
- removable mass storage media, 388
- replication, asynchronous/synchronous, 402
- repositioning servers, 130
- request domain identifiers (RDIs), 423
- Request for Comments.** *See* RFCs
- requests**
 - ARP (Address Resolution Protocol), 49
 - GET operations, 117
- requirements, 26**
 - NPV (N_Port Virtualization) connections, 476
- reserved IDs (VLANs), 84-85**
- resilience, applications, 37**
- resistors, 28**
- Resource Manager, 190**
- resources**
 - allocation, 17
 - VLANs (*virtual LANs*), 206
 - VRF (*Virtual Routing and Forwarding*), 101-103
 - classes, 155
 - Fibre Channel, 417
 - load balancers, 693
 - pooling, 794
 - sharing, 85
 - VDCs (virtual device contexts), 202-211
 - global*, 225
 - templates*, 211-213
 - virtual context (ACE) allocation, 145-156
- responses**
 - ARP (Address Resolution Protocol), 49
 - optimization, 130
 - server behavior, 118
- restarting fabrics, 424**
- results, LACP (Link Aggregation Control Protocol), 237**
- revisions, files, 391**
- RFCs (Request for Comments), 12**
- RHI (Route Health Injection), 775-777**
- RIB (Routing Information Base), 87, 92, 334**
- RIP (Routing Information Protocol)**
- RISC (Reduced Instruction Set Computing) servers, 561**
- risks, Layer 2 extensions, 325**
- RJ-45 connectors, 29**
- RL (rate limiter), 503**
- RLOC (Routing Locator), 779**
- Role-Based Access Control.** *See* RBAC
- roles**
 - ports, 73
 - VDCs (virtual device contexts), 223

- virtual contexts (ACE), 172-174
- vPC (virtual PortChannel) verification, 251
- root bridges, 62
- Root Port Selection phase, 68
- Root Selection phase, 67
- round-robin predictors, 116
- round-trip time. *See* RTT
- Route Health Injection. *See* RHI
- router-on-a-stick design, 55, 58
- routers
 - HSRP (Hot Standby Router Protocol), 11-12
 - virtual context (ACE) design, 156-158
- routes
 - static, 87
 - types, memory ranges, 210
- routing, 771-775
 - deployment, 58
 - frames, 421
 - GRE (generic routing encapsulation), 352-354
 - internal, Layer 2 extensions, 380-382
 - IVR (Inter-VSAN Routing), 453, 464-473
 - layers, 36
 - MC-LAG (Multi-Chassis Link Aggregation Group), 351
 - overview of, 87
 - PBR (Policy-Based Routing), 765
 - protocols, 92-98
 - tables
 - Layer 3 switches*, 88
 - MDS-CORE*, 442
 - VRF (Virtual Routing and Forwarding), 45, 89-91, 199
 - See also* VRF
- Routing Information Base. *See* RIB
- Routing Information Protocol. *See*

- RIP
- Routing Locator (RLOC), 779
- RPO (Recovery Point Objective), 321
- RSCN (Registered State Change Notification), 426, 464
- RSH (Remote Shell), 750
- RSS (Receive Side Scaling), 725
- RSTP (Rapid Spanning Tree Protocol), 72
- RTO (Recovery Time Objective), 321
- RTSP (Real-Time Streaming Protocol), 167
- RTT (round-trip time), 763
- rules, access (Cisco ASA 1000V), 761
- RUs (rack units), 572

S

- SaaS (Software as a Service), 795
- SAM (SCSI Architectural Model), 395
- Samba, 398
- SANs (storage area networks), 392, 409
 - consolidation, 447-450
 - extensions
 - FCoE (Fibre Channel over Ethernet)*, 545
 - with TE*, 462
 - FCIP (Fibre Channel over IP), 453-464
 - Fibre Channel, 410-420
 - islands, 430-432
 - IVR (Inter-VSAN Routing), 464-473
 - LAN management separation, 546-555
 - NPV (N-Port Virtualization), 473-490
 - virtualization, 407-408, 453
- SAN tab (UCS Manager), 591

- SAS (Serial Attached SCSI), 395
- SATA (Serial Advanced Technology Attachment), 397
- SATA Tunneling Protocol. *See* STP
- Satellite Discovery Protocol. *See* SDP
- Satellite Registration Protocol. *See* SRP
- SBCCS (Single-Byte Command Code Set), 397, 410, 456
- scalability, 15, 21, 25
 - ACE virtual contexts, 110
 - cloud computing, 803
 - Fabric Extenders, 288
 - load balancers, 111, 135
 - network services, 736
 - private VLANs, 80
 - service profiles, 582
 - STP (Spanning Tree Protocol), 326
 - virtualization, 17, 454, 494
 - VLANs/VRF, 45
 - VM-FEX (Virtual Machine Fabric Extender), 657
 - vPCs (virtual PortChannels), 231
- scaling applications, 113
- SCM (Supply Chain Management), 3
- scope
 - VLAN (virtual local-area network) configuration, 85
 - VRF (Virtual Routing and Forwarding), 99
 - VSANs (virtual storage-area networks), 445-447
- SCR (State Change Registration) frame, 427
- scripting, 584
- scrub policy, 611
- SCSI (Small Computer Systems Interface), 392-396, 410
- SCSI Architectural Model. *See* SAM
- SDNs (Software-Defined Networks), 800-801
- SDP (Satellite Discovery Protocol), 296
- SDR (Single Data Rate), 567
- seasonal workloads, 653-654
- secondary storage, 8, 388
- secret identities, 453
- Secure HTTP. *See* HTTPS
- Secure Shell. *See* SSH
- Secure Sockets Layer. *See* SSL
- security, 111
 - areas, 47
 - ASAs (Adaptive Security Appliances), 179
 - ASDM (Adaptive Security Device Manager), 754
 - content, 58
 - information, 735
 - IPSec (IP Security), 774
 - load balancers, 136-137
 - MAC (Media Access Control), 601
 - policies, 737
 - Cisco ASA 1000V configuration, 761*
 - formatting, 745-747*
 - verification, 752*
 - properties, 659
 - segmentation, 187
 - TLS (Transport Layer Security), 130
 - VSG (Virtual Security Gateway), 742-753
 - web, 129
 - zones, 225
- segmentation, 133
 - data centers, 103-104
 - Ethernets, 47
 - out-of-order (OOO), 148
 - security, 187
 - VXLANs (Virtual eXtensible Local Area Networks), 699
- segregation, private VLANs, 78-83

selection

- Designated Ports Selection phase, 68
- predictors, best practices, 117
- Root Port Selection phase, 68
- Root Selection phase, 67
- sites, 775-641

Selective Acknowledgment Options (RFC 2018), 763**Send Diagnostic command, 393****sending traffic to Cisco ASA 1000V, 758****sequences**

- Fibre Channel, 415-417
- number calculations, 133

Serial Advanced Technology Attachment (SATA), 397**Serial Attached SCSI. *See* SAS****Serial Storage Architecture. *See* SSA****Server Message Block. *See* SMB****Server NAT, 124-125****Server-Provided MAC Address. *See* SPMA****servers, 19, 111**

- access (Unified Fabric designs), 536-542
- architectures, 560-562
- assignment, 606-607
- blade, 488
- boot order, 604-605
- clusters, 324
- connections, 34
- DHCP (Dynamic Host Configuration Protocol), 585
- Direct Server Return, 123-124
- domains, 584-595
- evolution, 559-560
- farms, 115
 - configuration, 166*
 - round-robin predictors, 116*
- interfaces, 591

load balancers. *See* load balancers *configuration, 115****offloading, 130-134*****mainframes, 560-561****migration, 325, 625****models, access, 287-291****NIC (network interface card) teaming, 239****pooling, 649-653****provisioning, 581*****activation, 583-585******automation, 581******industrializing, 637-653*****real, configuration, 164-167****responses, 118****reverse proxies, 128-130****RISC (Reduced Instruction Set Computing), 561****service profiles, associating to, 612-619****traditional formats, 571****troubleshooting, 737****unified access, deployment, 509-523****virtualization, 775-781****x86, 562*****hardware evolution, 562-572******memory, 566-569******time of market, 581******UCS (Unified Computing System), 578-579******virtualization, 572-578*****Servers tab (UCS Manager), 591, 593****serviceability, 559****service-level agreements. *See* SLAs****services****cloud computing, 795****Fibre Channel, 420-421****measurements, 794**

networks, 735

applications, 111

data paths, 738-740

See also vPath

definitions, 735-738

vPath-enabled, 740-771

VSG (Virtual Security Gateway), 742-753

OpenStack, 802

policies, 170

profiles, 581

associating to servers, 612-619

boot order, 604-605

building, 588-623

identifying, 594

industrializing server provisioning, 637-653

installing operating systems, 620-623

maintenance policies, 606

network definitions, 599-608

operational policies, 608-612

policies, 626-633

seasonal workloads, 653-654

server assignment, 606-607

storage definitions, 595-598

templates, 640-649

UCS (Unified Computing System), 586-588

virtual interface placement, 602-603

vNICs (virtual NICs), 707-709

proxy servers, 129

virtual service nodes, 740, 751

VSM (Network Services Manager), 792

Session Initiation Protocol. *See* SIP

sessions

CSR 1000v, 772

Etheranalyzer, 225

monitoring, 203

SPAN, 205-206

SSH (Secure Shell), 747

sets, zones, 430

SFF (Small Form Factor) committee, 397

SFP+ (Enhanced Small Form Factor Pluggable), 32

shaping, traffic, 660

sharing

buses, 392

linecards, 191

operating systems, 10

ports, 73

processes, 189

resources, 85

segments, 69

subnets, 59

VLANs, 177

shielded twisted-pair (STP), 29

show interface command, 459

Shugart Associates, 392

SIMMs (Single In-Line Memory Modules), 568

Simple Network Management Protocol. *See* SNMP

simplification, 26, 43

simulation, 12

Single-Byte Command Code Set (SBCCS), 397

single-context switches, 510-518

Single Data Rate. *See* SDR

Single In-Line Memory Modules. *See* SIMMs

single-mode fiber (SMF), 29, 396

single-root I/O virtualization (SR-IOV), 706

SIP (Session Initiation Protocol), 167

- sites, 358
 - selections, 775-641
 - VLANs, 358
- sizes, blocks, 391
- SLAs (service-level agreements), 3
- sliding windows, calculations, 133
- slow-start algorithms, 116
- Small Computer Systems Interface. *See* SCSI
- Small Form Factor (SFF) committee, 397
- Smart Port Macros, 664
- SMB (Server Message Block), 398
- SMF (single-mode fiber), 29, 396
- SMP (Symmetric Multi Processing), 566
- SnapMirror, 403
- SNMP (Simple Network Management Protocol), 20, 85, 117
- sockets, 564
- SOF (start of frame), 417, 434
- Software as a Service. *See* SaaS
- Software-Defined Networks. *See* SDNs
- software modularization, 5
- solid-state drives. *See* SSDs
- space, buffers (PAUSE frames), 499
- Spanning Tree Protocol. *See* STP
- SPAN sessions, 205-206
- speed, convergence, 232
- SPMA (Server-Provided MAC Address), 508
- SQL (Structured Query Language), 398
- SRAM (static RAM), 566
- SRDF (Symmetrix Remote Data Facility), 403
- SR-IOV (single-root I/O virtualization), 706
- SRP (Satellite Registration Protocol), 296, 298
- SSA (Serial Storage Architecture), 396
- SSDs (solid-state drives), 389
- SSH (Secure Shell), 163
 - to active VSM, 674
 - VSG (Virtual Security Gateway), 747
- SSL (Secure Sockets Layer), 130-133
- StackWise, 241
- standardization of cloud computing, 789-792
- standard protocols, 392
- standards
 - See also* IEEE
 - AES (Advanced Encryption Standard), 328
 - ANSI/TIA-942, 40-41, 287
 - encryption, 133
 - Ethernet, 29, 33
 - Fibre Channel, 410
 - NIST (National Institute of Standards and Technology), 794
 - physical layers, 29
- standby nodes, 325
- start of frame (SOF), 417, 434
- star topologies, 41
- State Change Registration (SCR)
 - frame, 427
- stateless computing, 587, 625-626
- states, ports, 70-72
- static pinning, 301
- static RAM. *See* SRAM
- static routes, 87
- static vNICs, creating, 599
- status
 - Fabric Extender verification, 298, 313
 - IPSec (IP Security), 774
 - NPV (N_Port Virtualization), 477
 - STP (Spanning Tree Protocol) verification, 274

- VCs (virtual circuits), 339
- virtual service nodes, 751
- vPCs (virtual PortChannels), 253
- VPLS (Virtual Private LAN Service), 343-346
- stickiness tables, 116**
- storage, 387**
 - access, 391-399
 - block-based, 392-397*
 - files, 397-398*
 - interfaces, 495*
 - mainframe, 396-397*
 - records, 398-399*
 - AMS (Adaptable Modular Storage), 792
 - data center devices, 387-391
 - definitions, 8, 595-598
 - devices, 4, 18
 - disk arrays, 389-390
 - HDDs (hard disk drives), 388-389
 - libraries, 390-391
 - memory. *See* memory
 - physical-to-logical, 8
 - tape drives, 390-391
 - VDC (virtual device context) configuration, 519-523
 - virtualization, 399-408
- storage area networks. *See* SANs**
- Storage Reference Architecture, 792**
- STP (SATA Tunneling Protocol), 397**
- STP (shielded twisted-pair), 29**
- STP (Spanning Tree Protocol), 37**
 - convergence, 66
 - DCI (Data Center Interconnect), 326
 - design, 232
 - Fabric Extender interfaces, 299-301
 - FabricPath, 272-276, 330
 - instances, 74-78
 - link utilization, 232-234
 - optimization, 72-73
 - PortChannel configuration, 236
 - ports, states, 70-72
 - scalability, 326
 - status verification, 274
 - switches, 63
 - VDCs (virtual device contexts), 192
 - VLANs (virtual local-area networks), 61-78
 - vPCs (virtual PortChannels), 254-259
- straight-through topologies (Fabric Extender), 305-308**
- stress tests, virtual contexts (ACE), 150**
- Structured Query Language. *See* SQL structures**
 - metadata, 391
 - virtualization, 17
- subareas, 20**
 - ACE virtual contexts, 110
 - Fabric Extenders, 288
 - service profiles, 582
 - technologies, 15
 - virtualization, 454, 494
 - VLANs/VRF, 45
 - VM-FEX (Virtual Machine Fabric Extender), 657
 - vPCs (virtual PortChannels), 231
 - VSANs (virtual storage-area networks), 409
- subnets, VLAN association to IP (Internet Protocol), 56-58**
- suboptimal traffic, 137**
- subtypes, 15, 21**
 - ACE virtual contexts, 110
 - Fabric Extenders, 288
 - network services, 736
 - service profiles, 582
 - virtualization, 454, 494
 - VLANs/VRF, 45

- VM-FEX (Virtual Machine Fabric Extender), 657
- vPCs (virtual PortChannels), 231
- VSANs (virtual storage-area networks), 409
- Supply Chain Management.** *See* **SCM**
- SVIs (Switch Virtual Interfaces)**, 58, 145, 669, 792
- SW_ILS (Switch Internal Link Service)**, 425
- Switch1**
 - configuration, 196-197
 - tasks, 198
 - VDCs (virtual device contexts), 194
- switches**
 - access, 38
 - aggregation, 36
 - behavior, 48
 - blade, 474
 - broadcast frames, 61
 - Cisco Nexus 1000v
 - architecture*, 661-663
 - communication modes*, 663-664
 - deployment*, 666-683
 - loop avoidance*, 685
 - online migrations*, 693-697
 - port profiles*, 664-666
 - consolidation, 187
 - CSS (Content Services Switch), 139
 - DCI (Data Center Interconnect), 328
 - default path cost for, 64
 - DVS (Distributed Virtual Switch), 661
 - Ethernets, 48
 - Fabric Extenders, 296-299, 310
 - fabric initialization, 422
 - Fibre Channel, 506
 - interfaces
 - port IDs*, 65
 - RSTP (Rapid Spanning Tree Protocol)*, 72
 - Layer 3, 58, 88
 - Layer 4, 120-121
 - Layer 7, 113, 120-121
 - MDS (Multilayer Data Switch), 410, 424, 456
 - MPLS (Multiprotocol Label Switching), 47
 - MST (Multiple Spanning Tree), 76-77
 - N7K, 519, 524
 - Nexus 1000v
 - application acceleration*, 763-771
 - vPath configuration*, 748, 758
 - Nexus 7000, 92
 - OTV (Overlay Transport Virtualization)*, 354-377
 - port groups*, 195
 - NX-OS, 63, 188
 - out-of-band management, 143
 - port connections, 69
 - production, 142
 - STP (Spanning Tree Protocol), 63, 233
 - triangular topologies, 233
 - VDCs (virtual device contexts), 183-184, 214
 - See also* VDCs
 - VLANs (virtual local-area networks), 50
 - VPLS (Virtual Private LAN Service), 343-346
 - VSL (virtual switch link), 241
 - VSS (Virtual Switch System), 241
 - vSwitch (Virtual Switch), 559-661
- Switch Internal Link Service.** *See* **SW_ILS**

switchmode fex-fabric command, 297

Switch Virtual Interfaces. *See* SVIs

symmetric connection management, 122

Symmetric Multi Processing. *See* SMP

Symmetrix Remote Data Facility (SRDF), 403

synchronous replication, 402

System Manager, 189

T

tables

MAC (Media Access Control) addresses, 691

MDS-CORE routing, 442

routing, 87-88

stickiness, 116

Tag Protocol Identifier. *See* TPI

tags

Bus&Tag communication, 396

IEEE 802.1Q, 54

tape drives, 390-391, 403

targets, SCSI (Small Computer Systems Interface), 392

tasks

server provisioning, 583

Switch1, 198

taxonomies

defining, 14

virtualization, 15-17, 21

TCNs (Topology Change Notifications), 331

TCP (Transmission Control Protocol), 120

offloading servers, 130, 133-134

port redirection, 126

timeouts, 123

TCP Flow Optimization (TFO), 763

TCP/IP (Transmission Control Protocol/Internet Protocol), 11

TE (Traffic Engineering), 333, 340

SAN extension with, 462

VSANs (virtual storage-area networks), 443

teaming NICs (network interface cards), 238

technologies

areas, 15, 18-20, 21

ACE virtual contexts, 110

Fabric Extenders, 288

service profiles, 582

virtualization, 454, 494

VM-FEX (Virtual Machine Fabric Extender), 657

vPCs (virtual PortChannels), 231

VSANs (virtual storage-area networks), 409

classifying, 14-21

network services, 736

VLANs/VRF, 45

wireless, 27

Telecommunications Industry Association (TIA), 287

Telnet, 163

templates

OVF (Open Virtualization Format), 743

resources (VDCs), 211-213

service profiles, 640-649

terminals, 4

termination

connections, 133

SSL (Secure Sockets Layer), 132

terminators, 28

tertiary storage, 388

testing

- IP (Internet Protocol) connections, 200
- logins (VLAN commands), 86
- predictors, 117
- stress tests, virtual contexts (ACE), 150
- Test Unit Ready command, 393**
- TFO (TCP Flow Optimization), 763**
- TFTP (Trivial File Transfer Protocol), 750**
- Thicknet (10BASE5), 27**
- Thinnet (10GBASE2), 27**
- thin-provisioning, 405**
- three-layer campus design, 34**
- three-way handshakes, 133**
- TIA (Telecommunications Industry Association), 287**
- timelines**
 - Ethernet protocol, 33-34
 - virtualization, 12-13
- timeouts, TCP (Transmission Control Protocol), 123**
- timers, 51**
- time-sharing mechanisms, 10**
- Time-to-Live. *See* TTL**
- TLS (Transport Layer Security), 130**
- tools**
 - KVM, 620
 - MPC (Modular Policy CLI), 162
 - performance monitors, 111
 - Web Stress Tool, 151
- Top-of-Rack (ToR) designs, 39, 289-291**
- topologies, 25**
 - DCNM-SAN, 531
 - deployment, 36
 - Fabric Extender, 291, 305-315
 - dual-homed, 309-315*
 - straight-through, 305-308*
 - FCIP (Fibre Channel over IP), 457
 - Fibre Channel, 412-413
 - hub-and-spoke, 329
 - IVR (Inter-VSAN Routing), 466
 - logical, 213
 - networks, 34-41
 - layers, 35*
 - physical layouts, 39-40*
 - Nexus 1000v virtual networking, 741
 - NPV (N_Port Virtualization), 476
 - OTV (Overlay Transport Virtualization), 359
 - SCSI (Small Computer Systems Interface), 392
 - segmentation, 104
 - star, 41
 - triangular switch, 233
 - UCS (Unified Computing System), 592
 - virtual networking, 668
 - VPLS (Virtual Private LAN Service), 343
- Topology Change Notifications. *See* TCNs**
- ToR (Top-of-Rack) designs, 39, 44**
- TPI (Tag Protocol Identifier), 53**
- traditional formats, servers, 571**
- traffic**
 - analysis, 58
 - Cisco ASA 1000V, sending to, 758
 - direction, 35
 - Direct Server Return, 123-124
 - FHRPs (first-hop redundancy protocols), 380
 - isolation, 47, 78-83
 - load balancing, 43
 - localization, 382
 - management
 - allowing to virtual contexts (ACE), 162-163*

NPV (*N_Port Virtualization*),
 482-486
 patterns, 245
 PortChannel configuration, 238
 routing, 87
 shaping, 660
 suboptimal, 137
 Telnet, 163
 VSG (Virtual Security Gateway),
 sending to, 747
Traffic Engineering. *See* TE
transfer length, 394
**transitions, STP (Spanning Tree
 Protocol),** 70
**transit VSANs (virtual storage-area
 networks),** 472
translation, addresses, 124-127
**Transmission Control Protocol/
 Internet Protocol.** *See* TCP/IP
transparency, 12
transparent mode, 126-127
Transport Layer Security. *See* TLS
trends, cables, 34
triangular switch topologies, 233
Trivial File Transfer Protocol. *See*
 TFTP
tromboning, 326, 384, 775
troubleshooting
 interference, 28
 server virtualization, 737
 STP (Spanning Tree Protocol), 37
trunks
 ISL (Inter-Switch-Link), 56
 private VLANs, 81
 VLANs (virtual local-area networks),
 52-56
 VSANs (virtual storage-area net-
 works), 434-438
TTL (Time-to-Live), 61, 269
tunnels

FCIP (Fibre Channel over IP), 455,
 461
 GRE (generic routing encapsulation),
 352
 IPSec (IP Security), 774
 Q-in-Q configuration, 378
 STP (SATA Tunneling Protocol), 397
 VRRP (Virtual Router Redundancy
 Protocol), 461
twisted-pair cables, 28-29, 34
types, 15, 21
 of ACE virtual contexts, 110
 of Ethernet ports, 72
 of Fabric Extenders, 288
 of network services, 736
 of optical fiber cabling, 29
 of ports (Fibre Channel), 412-413
 of predictors, 116
 of routes, 210
 of service profiles, 582
 of twinax cables, 32
 of twisted-pair cables, 28
 of virtualization, 15, 454, 494
 of VLANs/VRF, 45
 of VM-FEXs (Virtual Machine
 Fabric Extenders), 657
 of vPCs (virtual PortChannels), 231
 of VSANs (virtual storage-area net-
 works), 409

U

UCS (Unified Computing System),
 581
 firmware policies, 633-637
 policies, 626-633
 service profiles, 586-588
 topologies, 592
 VM (virtual machine) locations, 722
 vNICs (virtual NICs), 707-709

vPCs (virtual PortChannels), 686
 x86 servers, 578-579

UCS Manager, 588
 access, 590
 migration, 625

UCS VMware Integration wizard, 711-716

UDIMMs (Unregistered DIMMs), 568-569

UDP (User Datagram Protocol), 120
 port redirection, 126
 virtual contexts (ACE), 164

ULP (upper-layer protocol), 412

unallocated VDCs, 191
See also VDCs

unicast responses, 49

unified access servers
 deployment, 509-523
 single-context switch configuration, 510-518
 storage VDC configuration, 519-523

Unified Computing System. *See* UCS

Unified Fabric, 495, 535-545

Uniform Resource Locators. *See* URLs

Universally Unique Identifier. *See* UUID

University of Manchester, 9

UNPAUSE frames, 499

Unregistered DIMMs. *See* UDIMMs

unshielded twisted-pair (UTP), 29

updating
 BIOS firmware, 636
 LSUs (link-state updates), 424
 virtual networking, 672

uplinks, 35, 591

upper-layer protocol. *See* ULP

URLs (Uniform Resource Locators), 116

User Datagram Protocol. *See* UDP

users
 account configuration, 175
 login configuration, 176

utilization
 CPUs (central processing units), 203
 STP (Spanning Tree Protocol), 232-234

UTP (unshielded twisted-pair), 29

UUID (Universally Unique Identifier), 594

V

values, default path cost for switches, 64

Vblock, 792

vCons (virtual network interface connections), 602

VCs (virtual circuits), 338-339, 343

vDCs (virtual data centers), 745, 781-783, 785-789, 788
See also cloud computing

VDCs (virtual device contexts), 183-184, 378
 configuration, 190-202
 data center zones, 225-227
 management, 190, 214-224
operations, 214-216
out-of-band, 217-222
process failures, 216-217
RBAC (Role-Based Access Control), 222-224

names, 198
 overview of, 188-190
 resources
allocation, 202-211
global, 225
templates, 211-213
 storage configuration, 519-523
 uses of, 187-188

- virtualization, extending, 184-186
- VEM (Virtual Ethernet Module), 661, 671, 739-740**
- VE_Port (Virtual E_Port), 506**
- verification**
 - context creation, 147
 - Fabric Extender status, 298
 - FLOGI (Fabric Login), 513
 - hardware, 190
 - MST (Multiple Spanning Tree), 202
 - policies, 752
 - PortChannels (Cisco Nexus 1000v switches), 687
 - roles, 251
 - stateless computing, 625-626
 - status (Fabric Extender), 313
 - STP (Spanning Tree Protocol) status, 274
 - VDC (virtual device context)
 - names, 198*
 - resource allocation, 212*
 - virtual interfaces, 690
 - VNMC (Virtual Network Manager Center) bindings, 761
 - VSG (Virtual Security Gateway), 750
- versions of CIMC (Cisco Integrated Management Controller), 634**
- VFC (Virtual Fibre Channel), 504**
 - interfaces, 512, 525
 - PortChannel configuration, 528-531
- VFI (Virtual Forwarding Instance), 342**
- VF_Port (Virtual F_Port), 506**
- vHBAs (virtual HBAs), 596, 604**
- VIF (Virtual Interface), 292**
- VIP (virtual IP), 116, 122**
 - class maps, 169
 - nodes, 324
- virtual circuits. *See* VCs**
- virtual contexts (ACE), 109, 139-178**
 - ACE (Application Control Engine)
 - classification, 110
 - application networking services, 111
 - configuration, 163-171
 - fault tolerance, 177-178
 - integrating, 156-161
 - load balancers, 111-134
 - management, controlling access, 171-176
 - multitenant data centers, 179-181
 - resource allocation, 145-156
 - VLANs, sharing, 177
 - Virtual Data Center. *See* vDC**
 - virtual data planes, 352**
 - virtual device contexts. *See* VDCs**
 - Virtual E_Port. *See* VE_Port**
 - Virtual Ethernet Module. *See* VEM**
 - Virtual eXtensible Local Area Networks. *See* VXLANs**
 - Virtual Fibre Channel. *See* VFC**
 - Virtual Forwarding Instance. *See* VFI**
 - Virtual F_Port. *See* VF_Port**
 - virtual HBAs. *See* vHBAs**
 - Virtual Interface (VIF), 292**
 - virtual interfaces**
 - placement, 602-603
 - verification, 690
 - virtual IP. *See* VIP**
- virtualization**
 - ACE (Application Control Engine)
 - classification, 110
 - benefits of networks, 42-44
 - classification, 14-21, 454, 494
 - clusters, 737
 - defining, 12
 - definitions, 1, 8
 - extending, 184-186
 - FCIP (Fibre Channel over IP), 453-464
 - file systems, 406-407

- IP (Internet Protocol) addresses, 668
- LUNs (logical unit numbers), 404-406
- mainframe, 10
- nesting, 199-202
- network services. *See* networks; services
- NPV (N Port Virtualization), 453, 473-490
- origins of, 8-13
- pooling, 116
- routing, 771-775
- SANs (storage area networks), 407-408, 453
- scalability, 17
- servers, 775-781
- storage, 399-408
- STP instances, 75
- taxonomies, 15-17, 21
- timelines, 12-13
- troubleshooting, 737
- UCS (Unified Computing System), 587
- VLANs. *See* VLANs
- VRF. *See* VRF
- x86 servers, 572-578
- Virtualized Multiservice Data Center.** *See* VMDC
- Virtual Link, 506**
- virtual local-area networks. *See* VLANs
- virtual logical unit numbers. *See* vLUNs
- Virtual Machine Fabric Extender.** *See* VM-FEX
- virtual machines. *See* VMs
- Virtual Media option, 604**
- virtual memory, 8-9
- virtual networking, 658-661
 - external connections, 684-688
 - NX-OS features, 688-693
 - topologies, 668
 - updating, 672
 - VM-FEX (Virtual Machine Fabric Extender), 705-707
 - High Performance mode, 723-730*
 - migrations, 720-722*
 - VXLANS (Virtual eXtensible Local Area Networks), 697-705
- virtual network interface connections. *See* vCons
- Virtual Network Link (VN-Link), 293**
- Virtual Network Manager Center.** *See* VNMC
- Virtual Network Tag (VNTag), 293**
- virtual NICs. *See* vNICs
- Virtual N_Port.** *See* VN_Port
- Virtual PortChannel (vPC), 43**
- Virtual PortChannel Host-Mode (vPC-HM), 686**
- virtual PortChannel Plus. *See* vPC+
- virtual PortChannels. *See* vPCs
- Virtual Private Clouds.** *See* VPCs
- Virtual Private LAN Service.** *See* VPLS
- Virtual Private Networks.** *See* VPNs
- Virtual Router Redundancy Protocol.** *See* VRRP
- Virtual Routing and Forwarding.** *See* VRF
- Virtual Security Gateway.** *See* VSG
- virtual service nodes, 740, 751
- virtual storage-area networks. *See* VSANs
- Virtual Supervisor Module.** *See* VSM
- Virtual Switch.** *See* vSwitch
- virtual switch link. *See* VSL
- Virtual Switch System.** *See* VSS
- virtual tape library. *See* VTL
- Virtual Wide Area Application Services.** *See* vWAAS

- vlan all command, 55
- VLAN ID (VLAN Identifier), 54
- VLAN Identifier. *See* VLAN ID
- VLANs (virtual local-area networks), 45
 - classification, 45
 - cloud computing, 786
 - configuration, 209
 - defining, 49-56
 - identifiers, 659
 - Layer 2 extensions, 377-379
 - Layer 3, 58-61
 - misconceptions about, 56-61
 - native, 84
 - overview of, 83-87
 - private VLANs, 78-83
 - reserved IDs, 84-85
 - scope configuration, 85
 - site VLANs, 358
 - STP (Spanning Tree Protocol), 61-78
 - trunks, 52-56
 - VDCs (virtual device contexts), 203
 - virtual context creation, 145
 - virtual contexts (ACE), 177
- vLUNs (virtual logical unit numbers), 787
- VMDC (Virtualized Multiservice Data Center), 792
- VM-FEX (Virtual Machine Fabric Extender), 589, 657, 705-707
 - classification, 657
 - deployment, 707-720
 - High Performance mode, 723-730
 - online migrations, 720-722
 - VM (virtual machine) migration, 716-720
 - vSphere (VMware) deployment, 709-710
- VM-Mobility, 779
- VM-Nomad, 770
- vMotion, 660
 - Long Distance, 693
 - port profiles, 682
- VMs (virtual machines), 8, 574-578
 - attributes, 751
 - characteristics, 690
 - cloud computing, 787
 - emulation, 659
 - IP (Internet Protocol) addresses on VSG, 752
 - names, 752
 - VM-FEX (Virtual Machine Fabric Extender) migration, 716-720
- VM tab (UCS Manager), 591
- VMware, 573-574
 - UCS VMware Integration wizard, 711-716
 - vSphere, deploying VM-FEX, 709-710
 - vSwitch (Virtual Switch), 659
- vNICs (virtual NICs), 599, 706, 771
 - iSCSI (Internet SCSI), 604
 - options, 604
 - UCS (Unified Computing System), 707-709
- VN-Link (Virtual Network Link), 293
- VNMC (Virtual Network Manager Center), 742-743
 - bindings, 760
 - Cisco ASA 1000V registration, 756
 - security policies, creating, 745-747
- VN_Port (Virtual N_Port), 506
- VNTag (Virtual Network Tag), 293
- volume detection, 440
- vPath, 738-740
 - connections, 768
 - network services, 740-771
 - Nexus 1000v switches, 748
 - WAAS (Wide Area Application Services), 766

- vPC+ (virtual PortChannel Plus), 276-280
- vPC-HM (Virtual PortChannel Host-Mode), 686
- vPCs (virtual PortChannels), 43, 231, 241-265, 466
 - classification, 231
 - configuration, 247-254
 - defining, 242-247
 - domains, defining, 248
 - FCoE (Fibre Channel over Ethernet), 538-540
 - first-hop routing protocols, 259-265
 - Layer 2
 - extensions*, 328-330
 - multipathing*, 265-280
 - link failures, 245
 - role verification, 251
 - status, 253
 - STP (Spanning Tree Protocol), 254-259
 - UCS (Unified Computing System), 686
- VPCs (Virtual Private Clouds), 797
- VPLS (Virtual Private LAN Service), 319, 333, 342-351, 803
- VPNs (Virtual Private Networks), 47, 333
- VRF (Virtual Routing and Forwarding), 45, 89
 - classification, 45
 - cloud computing, 787
 - configuration, 90-91, 199
 - Layer 2 extensions, 333
 - management planes, 98-101
 - resource allocation control, 101-103
 - routing protocols, 92-98
 - scope configuration, 99
 - VDCs (virtual device contexts), 203
 - VRF-awareness, 100-101
- vrf member interface command, 91
- VRRP (Virtual Router Redundancy Protocol), 12, 380, 460
- VSANs (virtual storage-area networks), 409
 - defining, 430-447
 - formatting, 432-434
 - FSPF (Fabric Shortest Path First) protocol, 442-444
 - IVR (Inter-VSAN Routing), 464-473
 - scoping, 445-447
 - transit, 472
 - trunking, 434-438
 - zones, 439-442
- VSG (Virtual Security Gateway), 735, 737, 742-753
 - firewall configuration, 746
 - installation, 743-745
 - management model, 751
- VSL (virtual switch link), 241
- VSM (Network Services Manager), 792
- VSM (Virtual Supervisor Module), 661
 - installation, 671
 - interfaces, 662
 - migration, 682
 - registration, 744
- vSphere (VMware), 621
- VSS (Virtual Switch System), 241, 329, 350
- vSwitch (Virtual Switch), 659
 - configuration, 673
 - overview of, 660-661
- VTL (virtual tape library), 403
- vWAAS (Virtual Wide Area Application Services), 735, 737
 - connectivity parameters, 765
 - optimized connections, 767
- VXLANs (Virtual eXtensible Local Area Networks), 697-705, 786
- vZones (virtual zones), 751, 754

W

WAAS (Wide Area Application Services), 764, 766

WANs (wide-area networks), 4, 87, 763

warm-standby sites, 322

wavelength-division multiplexing. *See* WDM

WCCP (Web Cache Control Protocol), 765

WDM (wavelength-division multiplexing), 327-328

Web Cache Control Protocol. *See* WCCP

web security, 129

websites

- load balancers, 111
- server farms for, 168

Web Stress Tool, 151

weight limitations, 118

Wide Area Application Services. *See* WAAS

wide-area networks. *See* WANs

Windows (Microsoft), 9

Windows Scaling (RFC 1323), 763

wireless technologies, 27

wizards

- Create Host Firmware Package, 634
- Create Service Profile (expert), 594, 626, 637
- Deploy OVF Template vCenter, 745
- Manage Hosts vCenter, 606
- UCS VMware Integration, 711-716

World Wide Names. *See* WWNs

Write command, 393

WWNs (World Wide Names), 414, 421, 486-488, 581

X

x86 servers, 562

- CPU evolution, 564-566

expansion busses, 569-571

- hardware evolution, 562-572
- memory, 566-569
- physical formats, 571-572
- time of market, 581

UCS (Unified Computing System), 578-579

- virtualization, 572-578

Xerox, 26

XML (Extensible Markup Language) plug-ins, 742

xTR, 779

Z

ZDA (Zone Distribution Area), 41

Zone Distribution Area (ZDA), 41

zones

- DMZs. *See* DMZs
- Fibre Channel, 429
- IVR (Inter-VSAN Routing), 467-472
- VDCs (virtual device contexts), 225-227
- virtual
 - VSANs (virtual storage-area networks), 439-442

Zone Server service, 429

zSeries mainframes, 397, 561