



NX-OS and Cisco Nexus Switching

Next-Generation Data Center Architectures
Second Edition

ciscopress.com

Ron Fuller, CCIE® No. 5851
David Jansen, CCIE® No. 5952
Matthew McPherson

FREE SAMPLE CHAPTER



SHARE WITH OTHERS

NX-OS and Cisco Nexus Switching

Next-Generation Data Center Architectures

Second Edition

Ron Fuller, David Jansen,
Matthew McPherson

Cisco Press

800 East 96th Street

Indianapolis, IN 46240

NX-OS and Cisco Nexus Switching

Next-Generation Data Center Architectures, Second Edition

Copyright © 2013 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.

Fourth Printing: June 2015

Library of Congress Cataloging-in-Publication data is on file.

ISBN-13: 978-1-58714-304-5

ISBN-10: 1-58714-304-6

Warning and Disclaimer

This book is designed to provide information about the Nexus Operating system and Nexus family of products. Every effort has been made to make this book as complete and as accurate as possible, but no warranty or fitness is implied.

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

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

Trademark Acknowledgments

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

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: Anand Sundaram
Associate Publisher: Dave Dusthimer	Executive Editor: Brett Bartow
Development Editor: Eleanor C. Bru	Copy Editor: Apostrophe Editing Services
Managing Editor: Sandra Schroeder	Technical Editors: Jeff Fry, Chad Hintz
Project Editor: Seth Kerney	Proofreader: Megan Wade
Editorial Assistant: Vanessa Evans	Indexer: Ken Johnson
Cover Designer: Gary Adair	Composition: Jake McFarland
Book Designer: Louisa Adair	



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, CCIP, CCNA, CCNP, CCSP, CCIW, 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, PowerPanels, ProConnect, ScriptShare, SenderBase, SMARTnet, Spectrum Expert, StackWise, The Fastest Way to Increase Your Internet Quotient, TransPath, WebEx, and the WebEx logo are registered trademarks of Cisco Systems, Inc. and/or its affiliates in the United States and certain other countries.

All other trademarks mentioned in this document or website are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (0812R)

About the Authors

Ron Fuller, CCIE No. 5851 (Routing and Switching/Storage Networking), is a technical marketing engineer (TME) on the Nexus 7000 team for Cisco. He has 21 years of experience in the industry and has held certifications from Novell, HP, Microsoft, ISC2, SNIA, and Cisco. His focus is working with customers worldwide to address their challenges with comprehensive end-to-end data center architectures and how they can best use Cisco technology to their advantage. He has had the opportunity to speak at Cisco Live on VDCs, NX-OS Multicast, and general design. He lives in Ohio with his wife and four wonderful children and enjoys travel and auto racing. He can be found on Twitter @ccie5851.

David Jansen, CCIE No. 5952, is a technical solutions architect for Data Center for Enterprise Central Area. David has more than 20 years' experience in the information technology industry. He has held multiple certifications from Microsoft, Novell, Checkpoint, and Cisco. His focus is to work with Enterprise customers to address end-to-end data center Enterprise architectures. David has been with Cisco for 15 years, has been working as a technical solutions architect for 6 years, and has provided unique experiences helping customers build architectures for Enterprise data centers. David has also been instrumental in developing and delivering the following Cisco Live presentations:

- Data Center Interconnect solutions to address L2 requirements between multiple data centers to meet application clusters and virtualization requirements (for the past five years)

- Next-generation data center architectures and operations tectorial and breakout sessions (for the past three years)

- Deploying NX-OS Nexus devices in the network infrastructure best practices

David holds a B.S.E. degree in computer science from the University of Michigan (Go Blue!) and an M.A. degree in adult education from Central Michigan University.

Matthew McPherson is a senior systems engineer and solutions architect for Cisco in the Central Select Operation, specializing in data center architectures. Matt has been with Cisco for more than 2 1/2 years and has more than 12 years of experience in the industry working for service providers and large enterprise customers in the financial and manufacturing verticals. He has held certifications from Juniper, Netscreen, and Cisco, and possesses a deep technical background in the areas of routing, switching, and security. His primary focus is working with strategic customers in greater Michigan to address their overall infrastructure challenges. He lives in Michigan with his wife and enjoys biking and collecting cars.

About the Technical Reviewers

Jeff Fry, CCIE No. 22061 (Routing and Switching), is a senior network engineer and data center architect for a large data center hosting company in the United States. Prior to working at the hosting company, Jeff designed and oversaw a globally diversified data center and call center network in a large call center environment. Jeff has designed and implemented several Cisco data center networks using the Catalyst 6500 and Nexus 7000, 5000, and 2000 devices. Jeff has more than 20 years of experience in the industry and has certifications from Microsoft, CompTIA, Novell, and Cisco. Jeff lives in Pennsylvania with his wife and their three sons, and they enjoy vacationing at Disney World. Jeff can be found on twitter at @fryguy_pa.

Chad Hintz, CCIE No. 15729, is a technical solutions architect for Cisco focusing on designing enterprise solutions for customers around Cisco data center technologies. He also holds three CCIEs in routing and switching, security, and storage. He has more than 10 years of experience in the industry and has held certifications from Novell, VMware, and Cisco. His focus is working with enterprise/commercial customers to address their challenges with comprehensive end-to-end data center architectures. Chad is a regular speaker at Cisco and industry conferences for data center technologies. Chad lives near Buffalo, New York, with his wife and two wonderful children.

Dedications

Ron Fuller: This book is dedicated to my loving wife Julie and my awesome children: Max, Sydney, Veronica, Emerson, and Wu. Thank you for showing me the world through your perspective and helping me appreciate the true joy of children. I can't thank you enough for believing in me when I told you I was going to write another book. Your support and encouragement has and always will be the key to any success I enjoy. Thank you for your love and support. I can't forget to acknowledge my parents, Gene and Joan. Your guidance and upbringing are what showed me what commitment and "getting it done right the first time" were all about. I've tried to make good use of my spare time and use my God-given gifts correctly.

David Jansen: This book is dedicated, again, to my loving wife Jenise and my three children: Kaitlyn, Joshua, and Jacob. You are the inspiration that gave me the dedication and determination to complete this project. Kaitlyn, Joshua, Jacob, you are three amazing kids; you are learning the skills to be the best at what you do and accomplish anything; keep up the great work. Thank you for all your love and support; I could not have completed (yet another) this without your help, support, and understanding. I also promise you that I am not going to take on another commitment like this for a long time; I promise. Second, I would also like to further dedicate this book to my parents, Michael and Dolores; you have given me proper tools, guidance, attitude, drive, and education to allow me to do what I do. I'm so grateful to God, who gives endurance, encouragement, and motivation to complete such a large project like this.

Matthew McPherson: This book is dedicated to my wonderful wife Carey who is the love of my life! Without your strength, support, and encouragement, I would never have been able to complete this project. I would also like to further dedicate this book to my parents, Steve and Kathy. Without your guidance and wisdom growing up, I would never have had the opportunity to write this book. Thank you both for always making me strive for greatness.

Acknowledgments

Ron Fuller: First, I'd like to thank my co-authors Dave Jansen and Matt McPherson. Dave, thank you for being such a good friend, a trusted co-worker, and a leader in our organization. You set the bar the rest of us try to reach. It has been great sharing a brain with you, and I look forward to more challenges and fun. Keep the goat rodeos coming! Matt, thank you for stepping in to help complete this project. You are awesome to work with, and your technical acumen is top-notch. People like you and Dave are the reason I love my job. We need to find some more concerts to go to, but hopefully I won't have to walk so far in the cold.

I'd like to thank Brett Bartow for his (almost) infinite patience with this project. It is a huge undertaking, and his persistence, understanding, and encouragement were greatly appreciated. You gave us the extra time we needed to make sure this was a fantastic and comprehensive second edition.

Chris Cleveland, it has been a pleasure working with you. Your guidance on the formatting and consistency makes the book something we all can be proud of. Thank you for making three propeller heads from Cisco look good.

To our technical editors, Jeff Fry and Chad Hintz: Thank you for the detail-oriented work and assistance making the book accurate and concise.

I'd like to thank my team, in particular Nikhil Kelshikar and Jeff Raymond, for giving me the time needed to complete this project. The rest of the TMEs, thank you for explaining some of these topics, some more than once, so I could make sure this book is accurate.

A special thanks to Lucien Avramov for the Python script example used in Chapter 7. I also appreciate the help getting around Moscow—my Russian isn't very good.

To my family, thank you for the time away from you while I worked on a book on a topic on things you don't get to see. Your understanding and support through the weekends and late nights are truly appreciated.

Finally, I want to thank God for the gifts he has given me and the opportunity to do what I love to do with people I enjoy to support my family. I couldn't ask for more.

David Jansen: This is my third book, and it has been a tremendous honor to work with the great people at Cisco Press. There are so many people to thank; I'm not sure where to begin. I'll start with Brett Bartow: Thank you for allowing me to publish another book; this is something I enjoy doing. As we are a year late on this, I appreciate your patience and tolerance on this project. I really appreciate you keeping me on track to complete the project in a timely manner, as we have missed several completion dates.

First, I would like to thank my friends and co-authors Ron Fuller and Matthew McPherson. Both of you guys are great friends and co-workers that ROCK. I can't think of two better people to work with to complete such a project. Cisco is one of the most amazing places I've ever worked, and it's people like you, who are wicked smart and a lot of fun to work with, that make it such a great place. I look forward to working on other projects in the future. I look forward to continuing to work with you and grow the

friendship into the future. We need to find another heavy concert to top the Megadeth one we went to together the winter of 2011; Brett, this is why we were a year late on the project ☺.

I want to thank Kevin Corbin for your contributions to the second editions. Thank you for being a great friend and peer co-worker.

Chris Cleveland, again it was a pleasure to work with you. Your expertise, professionalism, and follow-up as a development editor is unsurpassed; thank you for your hard work and quick turnaround; this helped to meet the deadlines set forth.

To our technical editors—Jeff Fry and Chad Hintz—thank you for the time, sharp eyes, and excellent comments/feedback; both of you created a ton more work for me just when I thought I was done. It was a pleasure having you as part of the team. I would like to thank Yuri Lukin for allowing me access to his lab to complete the Nexus 1000V chapter and all the services. Yuri you are a true professional; thank you for your time and expertise.

Thanks to my manager at Cisco, Robert Fisher—I appreciate your guidance and your trust in my ability to juggle the many work tasks along with extra projects like working on a book. Thank you for all your support Bob!

I would like to thank the heavy metal music world out there—it allowed me to stay focused when burning the midnight oil; I would not have been able to complete this without loud rock ‘n roll music. Thank you.

I want to thank my family for their support and understanding while I was working on this project late at night and being patient with me when my lack of rest may have made me a little less than pleasant to be around. I know it is also hard to sleep when Dad is downstairs writing and not realizing the db levels of the music while the rest of the family sleeps.

Most important, I would like to thank God for giving me the ability to complete such a task with dedication and determination and for providing me the skills, knowledge, and health needed to be successful in such a demanding profession.

Matthew McPherson: I'd like to thank both David and Ron for asking me to be a part of such an amazing opportunity. Dave, thank you again for being such an awesome friend and great mentor. You are a technical rock star who is always there to help in any situation, and for that reason I have the utmost respect for you. I cannot thank you enough your continued guidance and wisdom. Ron, you have become a great friend, and it's been a privilege to work more closely with you. I look forward to our continued friendship and am eagerly awaiting the next laser-sword match! It is so amazing to be at Cisco and have an opportunity to work with such amazing people!

Brett Bartow and Chris Cleveland, thank you both for your patience and support through my first publication. It has certainly been a privilege working with you both, and I look forward to starting on the third edition!

Chad Hintz and Jeff Fry, it has been great working with you both. Thank you for your time, attention to detail, and technical knowledge that have helped shape this book.

Gary Witt, thank you for stepping up to mentor me when I joined Cisco. You are such a great leader, and I certainly appreciate your insight and encouragement. I look forward to continuing our conversations and learning as much as I can.

I'd also like to thank my manager, Rajeev Grover. Your guidance and support for this project are truly appreciated!

Contents at a Glance

	Foreword	xxiii
	Introduction	xxiv
Chapter 1	Introduction to Cisco NX-OS	1
Chapter 2	Layer 2 Support and Configurations	59
Chapter 3	Layer 3 Support and Configurations	135
Chapter 4	IP Multicast Configuration	221
Chapter 5	Security	255
Chapter 6	High Availability	349
Chapter 7	Embedded Serviceability Features	385
Chapter 8	Unified Fabric	455
Chapter 9	Nexus 1000V	489
Chapter 10	Quality of Service (QoS)	643
Chapter 11	Overlay Transport Virtualization (OTV)	675
Chapter 12	Layer 3 Virtualization and Multiprotocol Label Switching (MPLS)	709
Chapter 13	LISP	729
Chapter 14	Nexus Migration Case Study	749
Index		789

Contents

Foreword xxiii

Introduction xxiv

Chapter 1 Introduction to Cisco NX-OS 1

NX-OS Overview 1

NX-OS Supported Platforms 3

NX-OS Licensing 7

Nexus 7000 7

Nexus 5500 8

Nexus 3000 8

Nexus 2000 9

Nexus 1000v 9

Installing the NX-OS License File 9

Cisco NX-OS and Cisco IOS Comparison 10

NX-OS User Modes 12

EXEC Command Mode 12

Global Configuration Command Mode 13

Interface Configuration Command Mode 13

Management Interfaces 14

Controller Processor (Supervisor Module) 15

Connectivity Management Processor (CMP) 16

Telnet 18

SSH 19

SNMP 23

DCNM 26

Managing System Files 28

File Systems 28

Configuration Files: Configuration Rollback 33

Operating System Files 35

Virtual Device Contexts 37

VDC Configuration 43

VDC Interface Allocation 46

Interface Allocation: N7K-M132XP-12 and L 46

Interface Allocation: N7K-F132XP-15 47

Interface Allocation: N7K-M108X2-12L 48

Interface Allocation: 10/100/1000 Modules 48

Interface Allocation on M2 Modules 52

Troubleshooting 54

show Commands 54

debug Commands 55

Topology 56

Further Reading 57

Chapter 2 Layer 2 Support and Configurations 59

Layer 2 Overview 59

Store-and-Forward Switching 60

Cut-Through Switching 60

Fabric Extension via the Nexus 2000 60

Configuring Nexus 2000 Using Static Pinning 61

Nexus 2000 Static Pinning Verification 62

Configuring Nexus 2000 Using Port-Channels 66

Nexus 2000 Static Pinning Verification 67

Layer 2 Forwarding on a Nexus 7000 69

L2 Forwarding Verification 70

VLANs 71

Configuring VLANs 72

VLAN Trunking Protocol 72

Assigning VLAN Membership 73

Verifying VLAN Configuration 74

Private VLANs 76

Configuring PVLANS 77

Verifying PVLAN Configuration 80

Spanning Tree Protocol 80

Rapid-PVST+ Configuration 82

Verifying Spanning Tree State for a VLAN 83

Spanning Tree Timers 84

MST Configuration 87

Additional Spanning Tree Configuration 91

Port Cost 91

Port Priority 94

Spanning Tree Toolkit 94

BPDUGuard 94

<i>BPDUFilter</i>	95
<i>RootGuard</i>	96
<i>LoopGuard</i>	97
<i>Dispute Mechanism</i>	98
<i>Bridge Assurance</i>	98
Spanning Tree Port Types	99
Virtualization Hosts	100
Configuring Layer 2 Interfaces	100
<i>Trunk Ports</i>	100
<i>Standard Host</i>	101
<i>Link to Virtualization Host</i>	101
<i>Port-Profiles</i>	102
Port-Channels	103
Assigning Physical Ports to a Port-Channel	104
Port-Channel Flow Control	107
Verifying Load Distribution Across a Port-Channel	108
Virtual Port-Channels	109
vPC Peer-Gateway	116
vPC Peer-Switch	116
ARP Synchronization	117
Unidirectional Link Detection	118
Cisco FabricPath	119
vPC+	127
Configuring vPC+	127
Summary	133
Chapter 3 Layer 3 Support and Configurations	135
EIGRP	135
EIGRP Operation	136
Configuring EIGRP	137
EIGRP Summarization	142
EIGRP Stub Routing	145
Securing EIGRP	147
EIGRP Redistribution	149
OSPF	154
OSPFv2 Configuration	154
OSPF Summarization	160

OSPF Stub Routing	163
Securing OSPF	167
OSPF Redistribution	169
OSPFv3 Configuration	177
IS-IS	178
IS-IS Configuration	178
BGP	183
BGP Configuration	184
BGP Neighbors	187
Securing BGP	190
BGP Peer Templates	192
Advertising BGP Networks	194
Modifying BGP Routing Metrics	197
Verifying BGP-Specific Configuration	198
First Hop Redundancy Protocols	198
HSRP	199
<i>HSRP Configuration</i>	199
<i>HSRP Priority and Preempt</i>	200
<i>Verifying the HSRP Configuration</i>	201
<i>Securing HSRP</i>	202
<i>HSRP Secondary Support</i>	204
<i>HSRP Support for IPv6</i>	204
VRRP	205
<i>VRRP Configuration</i>	205
<i>VRRP Priority and Preempt</i>	207
<i>Verifying VRRP Configuration</i>	208
<i>Securing VRRP</i>	208
<i>VRRP Secondary Support</i>	209
HSRP, VRRP, and vPC Interactions	210
GLBP	212
<i>GLBP Configuration</i>	212
<i>GLBP Priority and Preempt</i>	214
<i>Verifying GLBP Configuration</i>	214
<i>Securing GLBP</i>	215

	<i>GLBP Secondary Support</i>	218
	Summary	220
Chapter 4	IP Multicast Configuration	221
	Multicast Operation	221
	Multicast Distribution Trees	222
	Reverse Path Forwarding	225
	Protocol Independent Multicast (PIM)	225
	RPs	226
	PIM Configuration on Nexus 7000 and Nexus 5500	227
	Configuring Static RPs	230
	Configuring BSRs	232
	Configuring Auto-RP	235
	Configuring Anycast-RP	237
	Configuring SSM and Static RPF	239
	IGMP Operation	241
	IGMP Configuration on Nexus 7000	242
	IGMP Configuration on Nexus 5000	245
	IGMP Configuration on Nexus 1000V	246
	MSDP Configuration on Nexus 7000	248
	Administrative Scoping of Multicast RPs in PIM	250
	Configuring PIM Join and Prune Policies	252
	Multicast and Control Plane Policing (CoPP)	253
	Summary	253
Chapter 5	Security	255
	Configuring RADIUS	256
	RADIUS Configuration Distribution	259
	Configuring TACACS+	266
	Enabling TACACS+	266
	<i>TACACS+ Configuration Distribution</i>	267
	<i>Configuring the Global TACACS+ Keys</i>	268
	<i>Configuring the TACACS+ Server Hosts</i>	268
	<i>Configuring TACACS+ Server Groups</i>	269
	<i>Configuring TACACS+ Source Interface</i>	270
	Configuring SSH	275
	Cisco TrustSec	278
	Configuring AAA for Cisco TrustSec	281

- Defining Network Device Admission Control* 282
- Configuring the Nexus 7000 for 802.1x and SGA Features* 285
- SGT Assignment via ISE Server* 288
- Policy Component: IP to SGT Mapping* 290
- Policy Component: SGACL Creation* 292
- Configuring Cisco TrustSec: IEEE 802.1AE LinkSec 294
 - Layer 2 Solutions Between Data Centers 301
- Configuring IP ACLs 302
- Configuring MAC ACLs 305
- Configuring VLAN ACLs 307
- Configuring Port Security 308
 - Security Violations and Actions 311
- Configuring DHCP Snooping 313
- Configuring Dynamic ARP Inspection 316
 - Dynamic ARP Inspection Trust State 317
- Configuring IP Source Guard 321
- Configuring Keychain Management 322
- Configuring Traffic Storm Control 323
- Configuring Unicast RPF 325
- Configuring Control Plane Policing 327
- Configuring Rate Limits 335
- SNMPv3 340
- Summary 347

Chapter 6 High Availability 349

- Physical Redundancy 349
 - Redundant Power Supplies 350
 - Redundant Cooling System 352
 - Redundant Supervisors 355
 - Redundant Ethernet Out-of-Band (EOBC) 357
 - Redundant Fabric Modules 357
- Generic Online Diagnostics 358
 - Bootup Diagnostics 359
 - Runtime Diagnostics 360
 - On-Demand Diagnostics 365

NX-OS High-Availability Architecture	365
Process Modularity	366
Process Restart	368
Stateful Switchover	369
Nonstop Forwarding	370
In-Service Software Upgrades	370
Summary	383

Chapter 7 Embedded Serviceability Features 385

SPAN	386
SPAN on Nexus 7000	386
Configuring SPAN on Nexus 7000	387
SPAN on Nexus 5x00	392
Configuring SPAN on Nexus 5x00	393
SPAN on Nexus 1000V	397
Configuring SPAN on Nexus 1000V	398
ERSPAN on Nexus 1000V	400
ERSPAN on Nexus 7000	406
ERSPAN on Nexus 5x00	412
Embedded Analyzer	414
Smart Call Home	424
Smart Call Home Configuration	428
Configuration Checkpoint and Rollback on Nexus 7000	431
Checkpoint Creation and Rollback	432
Configuration Checkpoint and Rollback on Nexus 5x00	434
Checkpoint Creation and Rollback	435
NetFlow	437
Configuring NetFlow on Nexus 7000	438
Configuring NetFlow on Nexus 1000V	442
Network Time Protocol	444
Precision Time Protocol	445
IEEE 802.3az (Energy Efficient Ethernet)	447
Power On Auto-Provisioning	448
Python	449
Summary	454

Chapter 8 Unified Fabric 455

- Unified Fabric Overview 455
- Enabling Technologies 456
 - 10-Gigabit Ethernet 456
 - Fibre Channel over Ethernet 458
 - Single-Hop Fibre Channel over Ethernet 461
 - Multhop Fibre Channel over Ethernet 462
 - Storage VDC on Nexus 7000 463
- N-Port Virtualization 465
 - N-Port Identification Virtualization 466
 - FCoE NPV Mode 466
- Nexus 5x00 Unified Fabric Configuration 467
 - Single-Hop FCoE Configuration: Nexus 5x00 469
 - FCoE-NPV on Nexus 5x00 473
- Nexus 7000 Unified Fabric Configuration 477
- Summary 488

Chapter 9 Nexus 1000V 489

- Hypervisor and vSphere Introduction 489
- Nexus 1000V System Overview 490
- Nexus 1000V Switching Overview 494
- Nexus 1000V VSM Installation 496
 - Nexus 1000V Deployed on Nexus 1010 Virtual Services Blade 497
 - Registering the Nexus 1000V Plug-In to VMware Virtual Center Management Application 502*
 - Configuring the SVS Domain and Networking Characteristics 507*
 - Connecting the Nexus 1000V VSM to the vCenter Server 508*
 - Nexus 1000V Installation Management Center 510
 - VEM Installation Option on the Nexus 1000V Management Installation Center 519
 - vCenter Connection Option on the Nexus 1000V Management Installation Center 523
 - Creating the Uplink Profile 526
 - Adding the VEM to a ESX vSphere Host 528
 - Enabling the Telnet Server Process 536
 - Changing the VSM Hostname 536

Layer 3 Control	536
1000V Port Profiles	542
Virtual Network Management Center	552
Installing Virtual Network Management Center Software from OVA Downloaded from Cisco.com	553
Adding the VM-Manager for vCenter Connectivity in VNMC Management Application	564
Configuring the Cisco VNMC Policy-Agent on the 1000v VSM	570
Virtual Security Gateway	571
Install Virtual Security Gateway on the Nexus 1010	574
Configuring the Cisco VNMC Policy-Agent on the VSG	577
Verify That the VSG and VSM Are Registered Clients in VNMC	578
Creating a Tenant in VMMC	579
Virtual Extensible LAN	602
Deploying Virtual Extensible LAN	604
Nexus 1000v Network Analysis Module	629
Installing Nexus 1000v Network Analysis Module	630
Deploying the Nexus 1000v NAM as a Virtual Services Blade on the Nexus 1010	641
Summary	642
Chapter 10 Quality of Service (QoS)	643
QoS on Nexus 7000	646
Forwarding Architecture	646
Network-QoS Policies	648
Queuing Policies	650
QoS and Nexus 2000 Fabric Extenders	661
QoS and Nexus 7000 Virtual Device Contexts	663
QoS on Nexus 5x00	663
Forwarding Architecture	663
Network-QoS Policies	664
Queuing Policies	667
QoS and Nexus 2000 Fabric Extenders	668
QoS on Nexus 1000V	670
Forwarding Architecture	670
Classification in Nexus 1000V	670
Summary	674

Chapter 11 Overlay Transport Virtualization (OTV) 675

OTV Terminology and Concepts	677
OTV Control Plane	682
Multicast-Enabled Transport Infrastructure	687
Unicast-Enabled Transport Infrastructure	691
OTV Data-Plane	695
Data-Plane Multicast Traffic	697
OTV and QoS	698
Failure Isolation	698
STP Isolation	698
Unknown Unicast Handling with OTV	699
Broadcast Traffic Handling with OTV	699
Multihoming with OTV	700
OTV and ARP	700
First-Hop Routing Protocol Localization	702
Inbound Path Optimization	705
Summary	707

Chapter 12 Layer 3 Virtualization and Multiprotocol Label Switching (MPLS) 709

Virtual Routing and Forwarding	709
Predefined VRFs	710
VRF Operational Commands	713
VRF-Lite	713
MPLS Introduction	717
MPLS Terminology	718
LDP and Layer 3 VPNs	720
Quality of Service	723
Traffic Engineering	723
MPLS and IPv6: 6PE and 6VPE	725
Management and Troubleshooting	725
High Availability	725
Nexus Hardware Requirements and NX-OS Licensing for MPLS and VRF	726
Summary	727

Chapter 13 LISP 729

LISP Overview 729

LISP Terminology 730

LISP Prerequisites 731

LISP Control Plane 732

LISP Data Plane 733

Communicating Between LISP and non-LISP Sites 735

LISP Host Mobility with an Extended Subnet Mode 736

LISP Deployment Best Practices 746

Summary 746

Chapter 14 Nexus Migration Case Study 749

Existing Environment 749

Design Goals 750

The Design 751

Migration Plan 752

Premigration Steps 752

Maintenance Window #1 754

Maintenance Window #1 Summary 760

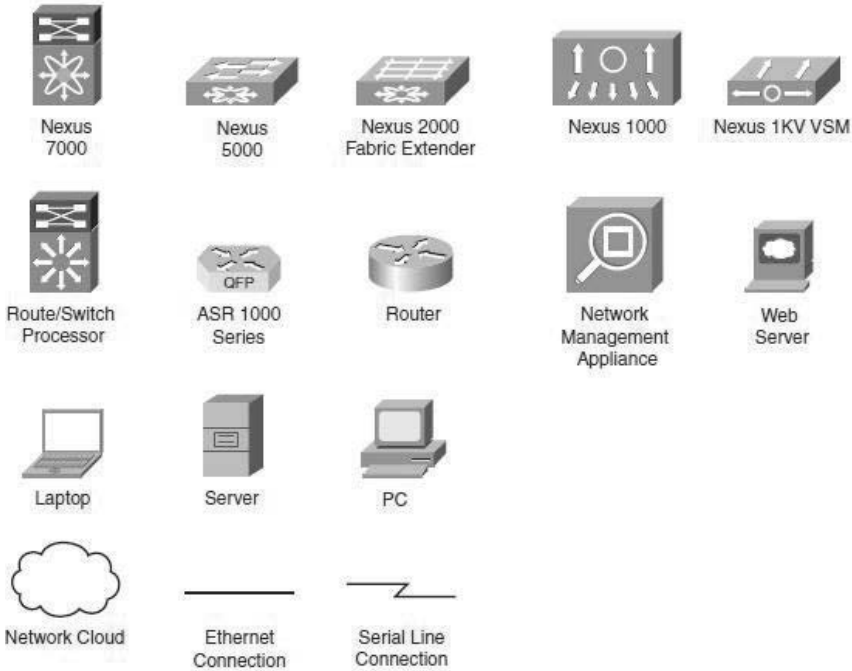
Maintenance Window #2 760

Ongoing Maintenance Windows 788

Summary 788

Index 789

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

Foreword

With more than 30,000 customers across every vertical, segment, and corner of the Earth, nobody can dispute that NX-OS is delivering on the requirements for the next-generation data center. In five quick years, seven product families have been released that are powered by NX-OS and provide end-to-end solutions in data centers large and small. These environments are as varied as ultra low latency trading environments, massively scalable data centers, cloud providers, and commercial and enterprise customer networks. NX-OS has the flexibility, reliability, security, and scalability to meet these demands and more.

The success of these products is driven by a team within Cisco that is committed to providing world-class solutions and solving challenges for customers with innovative technologies. Capabilities such as In-Service Software Upgrade (ISSU), modularity of the operating system, and stateful process restart lay a foundation for emerging technologies to build upon while preserving the investment in training and operations of the network. Game-changing capabilities such as Overlay Transport Virtualization (OTV), Locator Separator/ID Protocol (LISP), FabricPath, Fabric Extender architecture, vPath, Unified Ports, and dense 10G, 40G, and 100G interfaces provide customers a breadth of flexibility unparalleled in the industry—all running a common operating system, NX-OS.

To that end, a book like this can become a convenient reference for best practices deployment of these new technologies. It is written by two enterprise data center technology solutions architects and a technical marketing engineer on the Nexus 7000 team who all work with our customers on a daily basis and help them develop next-generation data center architectures. Their breadth of experience makes them perfect candidates to drive a project such as this.

We hope that as you read this book and learn more about the Nexus series of switches, and NX-OS specifically, you'll see the years of effort that made this product the Cisco flagship data center operating system now and in the years to come. Enjoy!

David Yen, SVP & GM
Data Center Business Unit
Cisco, San Jose

Introduction

The modern data center is rapidly changing and evolving to support the current and future demands of technology. At the center of this change is the network—the single entity that connects everything and touches all components of the data center. With that in mind, Cisco has launched a new series of switches, Nexus, based on a revolutionary new operating system, NX-OS, to meet these changes and provide a platform with the scalability, reliability, and comprehensive feature set required in the next-generation data center.

The purpose of this book is to provide a guide for the network administrator who might not be familiar with Nexus and NX-OS. It is intended to be used as a “go-to” resource for concise information on the most commonly used aspects of NX-OS across the Nexus 7000, 5000, 5500, and 1000V platforms.

Goals and Methods

The goal of this book is to provide best practice configurations to common internet-working scenarios involving Nexus products. Having been network administrators, the authors are conscious of the pressures and challenges with finding accurate and relevant information, especially on new technology. They intend this book to be a resource network administrators reach for first.

Although there might be more than one way to accomplish a networking requirement, this book focuses on the best way that minimizes operational complexity and maximizes supportability. The authors realize and respect that there might be corner-case scenarios that call for configurations not described in this book but sincerely hope they address the vast majority of common configurations.

Who Should Read This Book?

This book is targeted for the network administrator, consultant, or student looking for assistance with NX-OS configuration. It covers the three major Cisco Nexus products and highlights key features of them in a way that makes it easy to digest and implement.

How This Book Is Organized

This book has been organized following the OSI system model with the initial chapters starting with Layer 2 and then moving to Layer 3. Network-based services such as IP multicast, security, and high availability are then added. Next, the embedded serviceability features of NX-OS are explored, before moving to emerging data center architecture, Unified Fabric. With the drive toward virtualization, the need for increased visibility and control arises, and the Nexus 1000V meets these goals, which are covered next. Quality of service (QoS) is detailed before moving to the next topic, Overlay Transport Virtualization (OTV), where L2 segments can be safely extended between data centers.

The last chapter features a case study of an Enterprise customer who migrated from a Cisco Catalyst-based architecture to a Nexus-based one. The detailed step-by-step process is illustrated to provide a cookbook that can be used in many places.

Chapters 1 through 14 cover the following topics:

- **Chapter 1, “Introduction to Cisco NX-OS”:** Provides the reader with the foundation for building NX-OS configurations, including command-line interface (CLI) differences, virtualization capabilities, and basic file system management.
- **Chapter 2, “Layer 2 Support and Configurations”:** Focuses on the comprehensive suite of Layer 2 technologies supported by NX-OS, including vPC, Spanning Tree Protocol, and Cisco FabricPath.
- **Chapter 3, “Layer 3 Support and Configurations”:** Delves into the three most-common network Layer 3 protocols, including EIGRP, OSPF, and BGP. In addition, HSRP, GLBP, and VRRP are discussed.
- **Chapter 4, “IP Multicast Configuration”:** Provides the information needed to configure IP Multicast protocols such as PIM, Auto-RP, and MSDP.
- **Chapter 5, “Security”:** Focuses on the rich set of security protocols available in NX-OS, including CTS, SGTs, ACLs, CoPP, DAI, and more.
- **Chapter 6, “High Availability”:** Delves into the high-availability features built into NX-OS, including ISSU, stateful process restart, stateful switchover, and non-stop forwarding.
- **Chapter 7, “Embedded Serviceability Features”:** Provides the ability to leverage the embedded serviceability components in NX-OS, including SPAN, ERSPAN, configuration checkpoints and rollback, packet analysis, Smart Call Home, NTP, Python, and PoAP.
- **Chapter 8, “Unified Fabric”:** Explores the industry-leading capability for Nexus switches to unify storage and Ethernet fabrics with a focus on FCoE, NPV, and NPIV.
- **Chapter 9, “Nexus 1000V”:** Enables you to implement Nexus 1000V in a virtualized environment to maximum effect leveraging the VSM, VEM, and port profiles.
- **Chapter 10, “Quality of Service (QoS)”:** Illustrates the QoS capabilities of the Nexus platforms and covers the MQ CLI, queuing, and marking.
- **Chapter 11, “Overlay Transport Virtualization (OTV)”:** Delves into the details of this technology used to extend L2 networks across a L3 infrastructure.
- **Chapter 12, “Layer 3 Virtualization and Multiprotocol Label Switching (MPLS)”:** Covers how the integration of MPLS application components, including Layer 3 VPNs, traffic engineering, QoS, and mVPN-enable the development of highly efficient, scalable, and secure networks that guarantee service-level agreements.

- **Chapter 13, “LISP”:** Provides an introduction and overview of Locator ID Separation Protocol (LISP) in NX-OS and how this new routing architecture and paradigm shift decouples the server identity and the server location to allow for mobility, scalability, and security.
- **Chapter 14, “Nexus Migration Case Study”:** Detailed step-by-step description of a customer’s implementation of Nexus technology in their data center and the migration from a Catalyst-based architecture.

Unified Fabric

This chapter covers the following topics:

- Unified Fabric overview
- Enabling technologies
- Nexus 5x00 Unified Fabric configuration
- Nexus 7000 Unified Fabric configuration
- Cisco MDS Unified Fabric configuration

The Nexus family of switches represents a revolutionary approach to I/O within the data center referred to as Unified Fabric.

Unified Fabric Overview

One of the biggest trends in data centers today is consolidation, which can mean many different things. In some cases, consolidation refers to a physical consolidation of data centers where dozens or even hundreds of data centers are geographically dispersed and consolidated into a smaller number of large data centers. Consolidation can also exist within a data center where a large number of underutilized physical servers are consolidated, usually by leveraging some type of virtualization technology, into a smaller number of physical servers. Although virtualization offers many benefits, including consolidation of processors, memory, and storage, little is done to consolidate the amount of adapters, cables, and ports within the data center. In most virtualization implementations, there is actually a requirement for more adapters, cables, and ports to achieve the dense I/O requirements associated with virtualization. Data centers today contain multiple network fabrics that require discreet connectivity components to each fabric.

I/O consolidation is a trend within data centers that refers to the capability to aggregate connectivity to multiple fabrics into a single or redundant pair of adapters, cables, and

port. Although new technologies have emerged to enable this consolidation to occur, the concept is not new. Fibre Channel, iSCSI, Infiniband, and others were all introduced in an attempt to consolidate I/O. Although the merits or consolidation capabilities of each of these technologies might be open to debate, for one reason or another, all failed to reach mainstream adoption as the single fabric for all I/O requirements.

As a consolidation technology, Unified Fabric offers several benefits to customers, including

- **Lower capital expenditures:** Through the reduction of adapters, cables, and ports required within the infrastructure.
- **Lower operational expenses:** Through the reduction of adapters, cables, and ports drawing power within the data center.
- **Reduced deployment cycles:** Unified Fabric provides a wire-once model, in which all LAN, SAN, IPC, and management traffic is available to every server without requiring additional connectivity components.
- **Higher availability:** Quite simply, fewer adapters and ports means fewer components that could fail.

Enabling Technologies

Ethernet represents an ideal candidate for I/O consolidation. Ethernet is a well-understood and widely deployed medium that has taken on many consolidation efforts already. Ethernet has been used to consolidate other transport technologies such as FDDI, Token Ring, ATM, and Frame Relay networking technologies. It is agnostic from an upper layer perspective in that IP, IPX, AppleTalk, and others have used Ethernet as transport. More recently, Ethernet and IP have been used to consolidate voice and data networks. From a financial aspect, there is a tremendous investment in Ethernet that also must be taken into account.

For all the positive characteristics of Ethernet, there are several drawbacks of looking to Ethernet as an I/O consolidation technology. Ethernet has traditionally not been a lossless transport and relied on other protocols to guarantee delivery. In addition, a large portion of Ethernet networks range in speed from 100 Mbps to 1 Gbps and are not equipped to deal with the higher-bandwidth applications such as storage.

New hardware and technology standards are emerging that will enable Ethernet to overcome these limitations and become the leading candidate for consolidation.

10-Gigabit Ethernet

10-Gigabit Ethernet (10GbE) represents the next major speed transition for Ethernet technology. Like earlier transitions, 10GbE started as a technology reserved for backbone applications in the core of the network. New advances in optic and cabling technologies have made the price points for 10GbE attractive as a server access technology as well.

The desire for 10GbE as a server access technology is driven by advances in computer technology in the way of multisolet/multicore, larger memory capacity, and virtualization technology. In some cases, 10GbE is a requirement simply for the amount of network throughput required for a device. In other cases, however, the economics associated with multiple 1-G ports versus a single 10GbE port might drive the consolidation alone. In addition, 10GbE becoming the de facto standard for LAN-on-motherboard implementations is driving this adoption.

In addition to enabling higher transmission speeds, current 10GbE offerings provide a suite of extensions to traditional Ethernet. These extensions are standardized within IEEE 802.1 Data Center Bridging. Data Center Bridging is an umbrella referring to a collection of specific standards within IEEE 802.1, which are as follows:

- **Priority-based flow control (PFC; IEEE 802.1Qbb):** One of the basic challenges associated with I/O consolidation is that different protocols place different requirements on the underlying transport. IP traffic is designed to operate in large wide area network (WAN) environments that are global in scale, and as such applies mechanisms at higher layers to account for packet loss, for example, Transmission Control Protocol (TCP). Because of the capabilities of the upper layer protocols, underlying transports can experience packet loss and in some cases even require some loss to operate in the most efficient manner. Storage area networks (SANs), on the other hand, are typically smaller in scale than WAN environments. These protocols typically provide no guaranteed delivery mechanisms within the protocol and instead rely solely on the underlying transport to be completely lossless. Ethernet networks traditionally do not provide this lossless behavior for a number of reasons including collisions, link errors, or most commonly congestion. Congestion can be avoided with the implementation of *pause* frames. When a receiving node begins to experience congestion, it transmits a pause frame to the transmitting station, notifying it to stop sending frames for a period of time. Although this link-level pause creates a lossless link, it does so at the expense of performance for protocols equipped to deal with it in a more elegant manner. PFC solves this problem by enabling a pause frame to be sent only for a given Class of Service (CoS) value. This per-priority pause enables LAN and SAN traffic to coexist on a single link between two devices.
- **Enhanced transmission selection (ETS; IEEE 802.1Qaz):** The move to multiple 1-Gbps connections is done primarily for two reasons:
 - The aggregate throughput for a given connection exceeds 1 Gbps; this is straightforward but is not always the only reason that multiple 1-Gbps links are used.
 - To provide a separation of traffic, guaranteeing that one class of traffic will not interfere with the functionality of other classes. ETS provides a way to allocate bandwidth for each traffic class across a shared link. Each class of traffic can be guaranteed some portion of the link, and if a particular class doesn't use all the allocated bandwidth, that bandwidth can be shared with other classes.
- **Congestion notification (IEEE 802.1Qau):** Although PFC provides a mechanism for Ethernet to behave in a lossless manner, it is implemented on a hop-by-hop basis and

provides no way for multihop implementations. 802.1Qau is currently proposed as a mechanism to provide end-to-end congestion management. Through the use of backward congestion notification (BCN) and quantized congestion notification (QCN), Ethernet networks can provide dynamic rate limiting similar to what TCP provides only at Layer 2.

- **Data Center Bridging Capability Exchange Protocol extensions to LLDP (IEEE 802.1AB):** To negotiate the extensions to Ethernet on a specific connection and to ensure backward compatibility with legacy Ethernet networks, a negotiation protocol is required. Data Center Bridging Capability Exchange (DCBX) represents an extension to the industry standard Link Layer Discovery Protocol (LLDP). Using DCBX, two network devices can negotiate the support for PFC, ETS, and Congestion Management.

Fibre Channel over Ethernet

Fibre Channel over Ethernet (FCoE) represents the latest in standards-based I/O consolidation technologies. FCoE was approved within the FC-BB-5 working group of INCITS (formerly ANSI) T11. The beauty of FCoE is in its simplicity. As the name implies, FCoE is a mechanism that takes Fibre Channel (FC) frames and encapsulates them into an Ethernet. This simplicity enables for the existing skillsets and tools to be leveraged while reaping the benefits of a Unified I/O for LAN and SAN traffic.

FCoE provides two protocols to achieve Unified I/O:

- **FCoE:** The data plane protocol that encapsulates FC frames into an Ethernet header.
- **FCoE Initialization Protocol (FIP):** A control plane protocol that manages the login/logout process to the FC fabric.

Figure 8-1 provides a visual representation of FCoE.




Figure 8-1 *Fibre Channel over Ethernet*

When Fibre Channel frames are encapsulated in an Ethernet, the entire Fibre Channel frame, including the original Fibre Channel header, payload, and CRC are encapsulated in an Ethernet. Figure 8-2 depicts this.

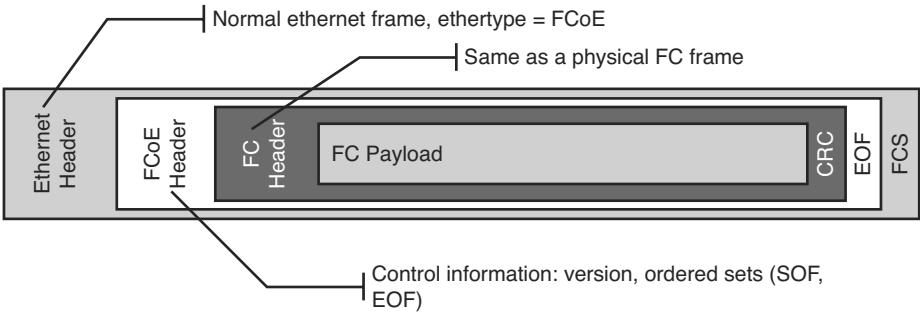


Figure 8-2 Fibre Channel Frame Encapsulated in an Ethernet

The ANSI T11 specifies the frame format for FCoE. It is a standard Ethernet frame with a new EtherType of 0x8906. Also note that the new Ethernet frame has a new Frame Check Sequence (FCS) created rather than using the FCS from the Fibre Channel frame. Figure 8-3 illustrates the FCoE frame format.

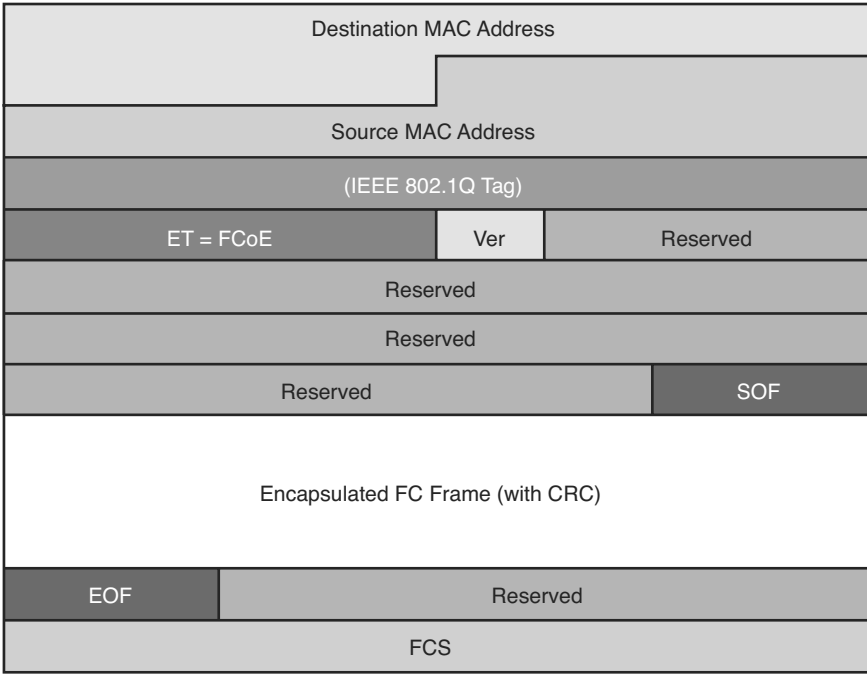


Figure 8-3 FCoE Frame Format

FCoE standards also define several new port types:

- **Virtual N_Port (VN_Port):** An N_Port that operates over an Ethernet link. N_Ports, also referred to as Node Ports, are the ports on hosts or storage arrays used to connect to the FC fabric.
- **Virtual F_Port (VF_Port):** An F_port that operates over an Ethernet link. F_Ports are switch or director ports that connect to a node.
- **Virtual E_Port (VE_Port):** An E_Port that operates over an Ethernet link. E_Ports or Expansion ports are used to connect Fibre Channel switches together; when two E_Ports are connected the link, it is an interswitch link (ISL).

To facilitate using FCoE an additional control plane protocol was needed and thus FCoE Initialization Protocol (FIP) was developed. FIP helps the FCoE perform VLAN discovery, assists the device in login (FLOGI) to the fabric, and finds key resources such as Fibre Channel Forwarders (FCFs). FIP is its own Ethertype (0x8914), which makes it easier to identify on a network and helps FIP Snooping devices identify FCoE traffic. Figure 8-4 depicts where FIP starts and ends and where FCoE takes over.

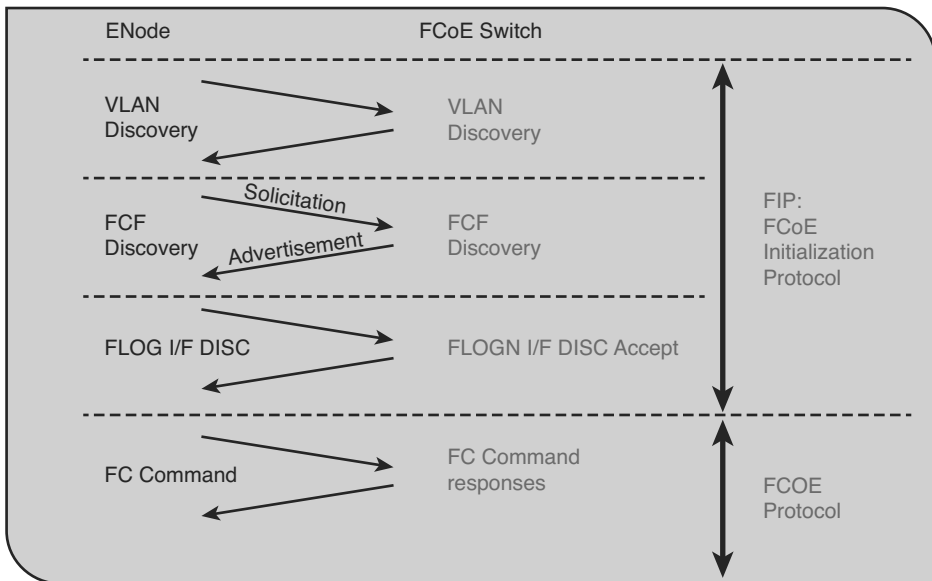


Figure 8-4 FIP Process

FIP can be leveraged by native FCoE-aware devices to help provide security against concerns such as spoofing MAC addresses of end nodes and helps simpler switches, such as FIP Snooping devices, learn about FCoE traffic. This awareness can provide security and QoS mechanisms that protect FCoE traffic from other Ethernet traffic and can help ensure a good experience with FCoE without the need to have a full FCoE stack on the switch. Currently the Nexus 4000 is the only Nexus device that supports FIP snooping.

Single-Hop Fibre Channel over Ethernet

Single-hop FCoE refers to an environment in which FCoE is enabled on one part of the network, frequently at the edge between the server and the directly connected network switch or fabric extender. In a single-hop topology the directly connected switch usually has native Fibre Channel ports which in turn uplink into an existing SAN, although you can have a complete network without any other fibre channel switches. Single-hop FCoE is the most commonly deployed FCoE model because of its double benefit of seamless interoperability into an existing SAN and the cost savings with a reduction in adapters, cabling, and optics to servers.

This reduction in cabling and adapters is accomplished through the use of a new adapter: Converged Network Adapter (CNA). CNAs have the capability to encapsulate Fibre Channel frames into Ethernet and use a 10GbE Ethernet interface to transmit both native Ethernet/IP traffic and storage traffic to the directly connected network switch or fabric extender. The CNA's drivers dictate how it appears to the underlying operating system, but in most cases it appears as a separate Ethernet card and separate Fibre Channel Host Bus Adapter (HBA).

Figure 8-5 shows how a CNA appears in Device Manager of a Microsoft Windows Server.



Figure 8-5 CNA in Device Manager

Using CNAs in a server, a typical single-hop FCoE topology would look like Figure 8-6 where a server is connected to Nexus 5x00 switches via Ethernet interfaces. The Nexus

5x00 switches have both Ethernet and native Fibre Channel interfaces for connectivity to the rest of the network topology. The fibre channel interfaces connect to native fibre channel ports on the Cisco MDS switches, and the Ethernet interfaces connect to the Ethernet interfaces on the Nexus 7000 switches. The FCoE traffic is transported only across the first or single hop from the server to the network switch. The current implementation of the Cisco Unified Computing System (UCS) uses single-hop FCoE between the UCS blade servers and the UCS Fabric Interconnects.




Figure 8-6 *Single-Hop FCoE Network Topology*

Multihop Fibre Channel over Ethernet

Building on the implementations of single-hop FCoE, multihop FCoE topologies can be created. As illustrated in Figure 8-6, native fibre channel links exist between the Nexus 5x00 and the Cisco MDS Fibre Channel switches, whereas separate Ethernet links interconnect the Nexus 5x00 and Nexus 7000. With multihop FCoE, topologies can be created where the native fibre channel links are not needed, and both fibre channel and Ethernet traffic use Ethernet interfaces.

The benefit of multihop FCoE is to simplify the topology and reduce the number of native fibre channel ports required in the network as a whole. Multihop FCoE takes the same principles of encapsulating fibre channel frames in Ethernet and uses it for switch-to-switch connections, referred to as Inter-Switch Links (ISL) in the Fibre Channel world, and uses the VE port capability in the switches.

Figure 8-7 shows a multihop FCoE topology where the server connects via CNAs to Nexus 5x00s, which in turn connect to Nexus 7000 series switches via the Ethernet carrying FCoE. The storage array is directly connected to the Nexus 7000 via FCoE as well.

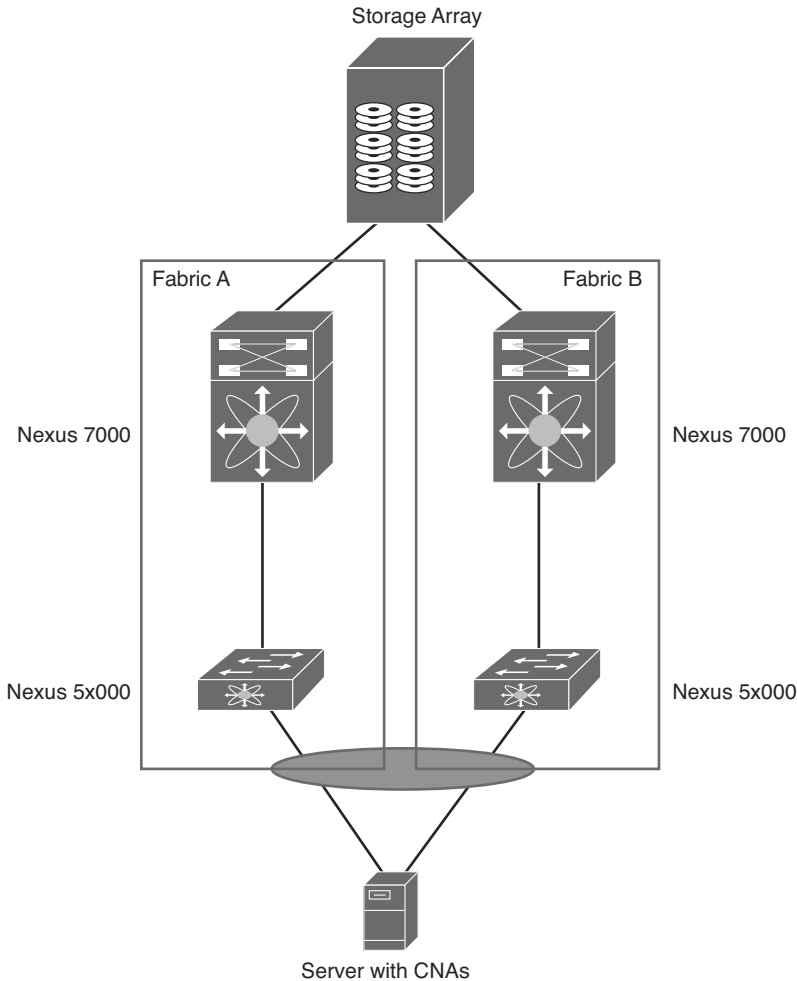


Figure 8-7 *Multihop FCoE Topology*

Storage VDC on Nexus 7000

One of the building blocks in a multihop FCoE topology is the storage Virtual Device Context (VDC) on the Nexus 7000. VDCs are discussed in detail in Chapter 1, “Introduction to Cisco NX-OS,” and the focus in this chapter is on the Storage VDC and its use in a multihop FCoE topology. VDC is a capability of the Nexus 7000 series switches that enables a network administrator to logically virtualize the Nexus 7000 into

multiple logical devices. The storage VDC is a special VDC that enables the virtualization of storage resources on the switch. This enables in essence a “virtual MDS” inside the Nexus 7000 that participates fully in the FCoE network as a full fibre channel forwarder (FCF).

With a Storage VDC network, administrators can provide the storage team a context that allows the storage team to manage their own interfaces; configurations; and fibre channel-specific attributes such as zones, zonesets, and aliases. Figure 8-8 shows how a storage VDC can be implanted in an existing topology where single-hop FCoE was initially deployed and then multihop FCoE was added. The storage VDC was created with VE ports connecting downstream to the Nexus 7000 and VE port to the Cisco MDS fibre channel director.




Figure 8-8 Storage VDC on the Nexus 7000

The storage VDC has some requirements that are unique to this type of VDC as storage traffic is traversing it. The first requirement is that the storage VDC can support only interfaces hosted on the F1 or F2/F2e series of modules. These modules support the capability to provide lossless Ethernet and as such are only suitable for doing FCoE. The VDC allocation process in NX-OS does not allow for other types of modules to have interfaces in a VDC that has been defined as a storage VDC.

Note To use FCoE with the F2/F2e series module, the switch must have a Supervisor 2 or Supervisor 2E installed. FCoE on F2/F2e is not supported with the Supervisor 1 module.

In addition to requiring F1 or F2/F2e series modules, the storage VDC cannot run non-storage related protocols. You cannot enable features such as OSPF, vPC, PIM, or other Ethernet/IP protocols in the storage VDC. The only features allowed are directly related to storage. Finally, the default VDC cannot be configured as a storage VDC.

N-Port Virtualization

The fibre channel module of the Nexus 5x00 series switch can operate in two modes:

- Fabric
- NPV (N-Port Virtualization)

When in fabric mode, the switch module operates as any switch in a fibre channel network does.

Fabric mode switches have the following characteristics:

- Unique domain ID per virtual storage area network (VSAN)
- Participation in all domain services (zoning, fabric security, Fibre Channel Identification [FCID] allocation, and so on)
- Support for interoperability modes

When the fibre channel module is configured in NPV mode, it does not operate as a typical fibre channel switch; instead leveraging a service, NPIV, on the upstream or core fibre channel switch for domain services. The switch operates in a similar fashion as an NPIV-enabled host on the fabric. The advantage NPV provides the network administrator is the control of domain IDs and points of management on a fibre channel network as it scales.

Note The fibre channel specification supports 239 domain IDs per VSAN; however, the reality is that many SAN vendors recommend and support a much lower number. Consult your storage vendor (Original Storage Manufacturer [OSM]) for specific scalability numbers.

Additional benefits of NPV include the capability to manage the fibre channel switch as a discrete entity for tasks such as software management and debugging the fibre channel network. NPV also enables network administrators to connect FCoE hosts to non-FCoE-enabled SANs and simplifies third-party interoperability concerns because the NPV enabled fibre channel module does not participate in domain operations or perform local switching. This enables multivendor topologies to be implemented without the restrictions the interoperability mode requires.

The fibre channel module in the Nexus 5x00 creates a new port type to the fibre channel network when in NPV mode: the NP-port. The NP-port proxies fabric login (FLOGI) requests from end stations and converts them to Fabric Discoveries (FDISC) dynamically and transparently to the end device. The result is that end systems see the NPV-enabled switch as a Fabric Port (F-port) and the upstream/core switch sees the NPV-enabled switch as an F-port as well. Figure 8-9 illustrates the port roles used in an NPV-enabled network.




Figure 8-9 Port Roles in an NPV-Enabled Network

Note Enabling NPV mode can cause the current configuration to be erased and the device rebooted. It is therefore recommended that NPV be enabled prior to completing any additional configuration.

N-Port Identification Virtualization

A key component to enable the proper operation of NPV is the need for N-Port Identification Virtualization (NPIV) on the core/upstream fibre channel switch. NPIV is an industry-standard technology defined by the T11 committee as part of the Fibre Channel Link Services (FCoE) specification and enables multiple N Port IDs or FCIDs to share a single physical N Port. Prior to NPIV, it was not possible to have a system that used multiple logins per physical port—it was a one-login-to-one-port mapping. With the increasing adoption of technologies such as virtualization, the need to allow multiple logins was created. NPIV operates by using Fabric Discovery (FDISC) requests to obtain additional FCIDs.

FCoE NPV Mode

Building on Fibre Channel NPV mode, the Nexus 5x00 supports running in FCoE-NPV mode as well. FCoE-NPV brings similar benefits as the Fibre Channel NPV mode to a pure FCoE implementation. The switch still uses FIP snooping to determine FCoE traffic and to maintain separation and provide security with the benefits of minimized domain sprawl, simplified management, and fewer FCoE devices to manage. FCoE NPV also creates a new port type for the VNP (Virtual NPV Port). Figure 8-10 illustrates where the VNP port resides in an FCoE NPV topology.


```

FC_FEATURES_PKG          Yes - Unused Never -
VMFEX_FEATURE_PKG       No - Unused -
ENHANCED_LAYER2_PKG     No - Unused -
-----
N5K-1#

```

Example 8-2 shows how to enable the FCoE feature.

Example 8-2 Enabling FCoE

```

N5K-1# config
Enter configuration commands, one per line. End with CNTL/Z.
N5K-1(config)# feature fcoe
FC license checked out successfully
fc_plugin extracted successfully
FC plugin loaded successfully
FCoE manager enabled successfully
N5K-1(config)#
N5K-1(config)# show license usage
Feature                Ins  Lic  Status Expiry Date Comments
                        Count
-----
FCOE_NPV_PKG          No   -   Unused -
FM_SERVER_PKG         No   -   Unused -
ENTERPRISE_PKG        Yes  -   Unused Never -
FC_FEATURES_PKG       Yes  -   In use Never -
VMFEX_FEATURE_PKG     No   -   Unused -
ENHANCED_LAYER2_PKG   No   -   Unused -
-----
N5K-1(config)#

```

Enabling NPV mode requires a write erase and reboot, as demonstrated in Example 8-3.

Example 8-3 Enabling NPV Mode

```

N5K-1# config
Enter configuration commands, one per line. End with CNTL/Z.
N5K-1(config)# show license usage
Feature                Ins  Lic  Status Expiry Date Comments
                        Count
-----
FCOE_NPV_PKG          No   -   Unused -
FM_SERVER_PKG         No   -   Unused -
ENTERPRISE_PKG        Yes  -   Unused Never -
FC_FEATURES_PKG       Yes  -   In use Never -

```

```

VMFEX_FEATURE_PKG           No - Unused -
ENHANCED_LAYER2_PKG        No - Unused -
-----

N5K-1(config)# feature npv
Verify that boot variables are set and the changes are saved.
Changing to npv mode erases the current configuration and reboots the
switch in npv mode. Do you want to continue? (y/n):y
Shutdown Ports..
  writing reset reason 90,
2012 Jul 30 00:32:39 N5K-1 %$ VDC-1 %$ Jul 30 00:32:39 %KERN-0-
SYSTEM_MSG: Shutdown Ports.. - kernel
2012 Jul 30 00:32:39 N5K-1 %$ VDC-1 %$ Jul 30 00:32:39 %KERN-0-
SYSTEM_MSG: writINIT: Sending processes the TERM signal
Sending all processes the TERM signal...
Sending all processes the KILL signal...
Unmounting filesystems...
Restarting system.

```

Single-Hop FCoE Configuration: Nexus 5x00

Now that the switches are configured for FCoE and have NPV configured, the next step is to configure the interconnection between the upstream Fibre Channel switch and the Nexus 5x00. In this example, a Nexus 5010 is connected to a Cisco MDS 9500 Fibre Channel directory via a 4-Gb native Fibre Channel port.

The first step is to configure the MDS to use NPIV, configure the port, and add it to the correct VSAN. This enables the MDS to support multiple FLOGI on a physical interface (NPIV), and for good documentation a description is added to the physical interface before being enabled. Finally, the port is added to the correct VSAN, 10 in this example. Figure 8-11 shows the topology for this environment.




Figure 8-11 *Single-Hop FCoE with Nexus 5x00*

Example 8-4 shows how to configure the ISL between the MDS and the Nexus 5000.

Example 8-4 *Configuring the MDS Port*

```

CMHLAB-DC1-MDS1# config
CMHLAB-DC1-MDS1(config)# feature npiv
CMHLAB-DC1-MDS1(config)# interface fc3/4
CMHLAB-DC1-MDS1(config)# switchport description Connection to CMHLAB-DC1-TOR1 2/1
CMHLAB-DC1-MDS1(config)# switchport trunk mode off
CMHLAB-DC1-MDS1(config)# no shutdown
CMHLAB-DC1-MDS1(config)# vsan database
CMHLAB-DC1-MDS1(config-vsan-db)# vsan 10 interface fc3/4
CMHLAB-DC1-MDS1(config)# end
CMHLAB-DC1-MDS1#
CMHLAB-DC1-MDS1# show vsan membership interface fc3/4
fc3/4
    vsan:10
    allowed list:1-4078,4080-4093
CMHLAB-DC1-MDS1#

```

Next, the Nexus 5x00 needs to have a port configured for the connection to the MDS. The port is configured for the NP mode and added to the appropriate VSAN, 10 to match with the MDS configuration.

Example 8-5 shows how to configure the fibre channel uplink to the SAN core.

Example 8-5 *Configuring FC Uplink*

```

CMHLAB-DC1-TOR1# config
Enter configuration commands, one per line. End with CNTL/Z.
CMHLAB-DC1-TOR1(config)# int fc2/1
CMHLAB-DC1-TOR1(config-if)# switchport mode NP
CMHLAB-DC1-TOR1(config-if)# switchport description Connection to CMHLAB-DC1-MDS1
fc3/4
CMHLAB-DC1-TOR1(config-if)# no shutdown
CMHLAB-DC1-TOR1(config-if)# end
CMHLAB-DC1-TOR1#

CMHLAB-DC1-TOR1# show int fc2/1
fc2/1 is up
    Port description is Connection to CMHLAB-DC1-MDS1 fc3/4
    Hardware is Fibre Channel, SFP is short wave laser w/o OFC (SN)
    Port WWN is 20:41:00:0d:ec:a3:0d:00
    Admin port mode is NP, trunk mode is off
    snmp link state traps are enabled
    Port mode is NP
    Port vsan is 10
    Speed is 4 Gbps

```

```

Transmit B2B Credit is 16
Receive B2B Credit is 16
Receive data field Size is 2112
Beacon is turned off
1 minute input rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
1 minute output rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
  10055 frames input, 5625012 bytes
    0 discards, 0 errors
    0 CRC, 0 unknown class
    0 too long, 0 too short
  10054 frames output, 523260 bytes
    0 discards, 0 errors
  1 input OLS, 1 LRR, 0 NOS, 0 loop inits
  1 output OLS, 1 LRR, 0 NOS, 0 loop inits
last clearing of "show interface" counters never
  16 receive B2B credit remaining
  16 transmit B2B credit remaining
  0 low priority transmit B2B credit remaining
Interface last changed at Mon May 21 20:09:15 2012

```

```
CMHLAB-DC1-TOR1# show npv sta
```

```
npiv is enabled
```

```
disruptive load balancing is disabled
```

```
External Interfaces:
```

```
=====
```

```
Interface: fc2/1, VSAN: 10, FCID: 0x7c0020, State: Up
```

```
Number of External Interfaces: 1
```

```
Server Interfaces:
```

```
=====
```

```
Number of Server Interfaces: 0
```

```
CMHLAB-DC1-TOR1#
```

After the connection between the MDS and Nexus 5x00 is configured, the next task is to configure the FCoE VLAN to VSAN mapping, configure the Ethernet interface that connects to the server, and finally configure the Virtual Fibre Channel (VFC) interface. This process is shown in Example 8-6 and Example 8-7.

Example 8-6 *Configuring FCoE VLAN to VSAN Mapping*

```

CMHLAB-DC1-TOR1# config
Enter configuration commands, one per line. End with CNTL/Z.
CMHLAB-DC1-TOR1(config)# vlan 10
CMHLAB-DC1-TOR1(config-vlan)# fcoe vsan 10
CMHLAB-DC1-TOR1(config-vlan)# name FCOE-FabA
CMHLAB-DC1-TOR1(config-vlan)# end
CMHLAB-DC1-TOR1# show vlan fcoe

```

Original VLAN ID	Translated VSAN ID	Association State
10	10	Operational

```

CMHLAB-DC1-TOR1#

```

After the FCoE VLAN is configured and mapped to a fibre channel VSAN, the Ethernet port that connects to the server should be configured (refer to Example 8-7).

Example 8-7 *Configuring the Physical and VFC Interface for FCoE*

```

CMHLAB-DC1-TOR1# config
Enter configuration commands, one per line. End with CNTL/Z.
CMHLAB-DC1-TOR1(config)# interface Ethernet1/7
CMHLAB-DC1-TOR1(config-if)# description Connection to DEMOLAB-VM1 - Emulex CNA
CMHLAB-DC1-TOR1(config-if)# switchport mode trunk
CMHLAB-DC1-TOR1(config-if)# switchport trunk allowed vlan 10,101,301,401,701,801
CMHLAB-DC1-TOR1(config-if)# interface vfc17
CMHLAB-DC1-TOR1(config-if)# bind interface Ethernet1/7
CMHLAB-DC1-TOR1(config-if)# switchport description FCoE Interface for DEMOLAB-VM1
CMHLAB-DC1-TOR1(config-if)# no shutdown
CMHLAB-DC1-TOR1(config-if)# end
CMHLAB-DC1-TOR1# CMHLAB-DC1-TOR1# show int e1/7 trunk

```

Port	Native Vlan	Status	Port Channel
Eth1/7	1	trunking	--

```

-----
Port                Vlans Allowed on Trunk
-----
Eth1/7              10,101,301,401,701,801

```

```

-----
Port          Vlans Err-disabled on Trunk
-----
Eth1/7        none
-----

Port          STP Forwarding
-----
Eth1/7        10,101,301,401,701,801
-----

Port          Vlans in spanning tree forwarding state and not pruned
-----
Eth1/7        --
-----

Port          Vlans Forwarding on FabricPath
-----

```

```
CMHLAB-DC1-TOR1# show int vfc17
```

```
vfc17 is up
  Bound interface is Ethernet1/7
  Port description is FCoE Interface for DEMOLAB-VM1
  Hardware is Ethernet
  Port WWN is 20:10:00:0d:ec:a3:0d:3f
  Admin port mode is F, trunk mode is on
  snmp link state traps are enabled
  Port vsan is 10
  1 minute input rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
  1 minute output rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
    0 frames input, 0 bytes
    0 discards, 0 errors
    0 frames output, 0 bytes
    0 discards, 0 errors
  last clearing of "show interface" counters never
```

```
CMHLAB-DC1-TOR1#
```

FCoE-NPV on Nexus 5x00

Configuration of the FCoE NPV mode on a Nexus 5x00 switch is similar to the configuration for the Fibre Channel NPV mode. The main difference is the configuration of an Ethernet port for the ISL and the VNP port. Figure 8-12 shows the topology used for the FCoE-NPV examples.




Figure 8-12 FCoE NPV Configuration Between a Nexus 5000 and Nexus 7000

First, the FCoE NPV feature must be enabled, as shown in Example 8-8.

Note FCoE-NPV cannot be enabled if FCoE is already enabled; otherwise, the following message displays: ERROR: Cannot enable feature fcoe-npv because feature fcoe is enabled. Disable feature fcoe, reload the system, and try again.

Example 8-8 FCOE-NPV Feature Installation

```
N5K-1# config
Enter configuration commands, one per line. End with CNTL/Z.
N5K-1(config)# feature fcoe-npv
FCoE NPV license checked out successfully
fc_plugin extracted successfully
FC plugin loaded successfully
FCoE manager enabled successfully
FCoE NPV enabled on all modules successfully
N5K-1(config)# end
N5K-1#
```

After the feature is installed, the switch needs to be configured for the VSAN and VLAN mapping to associate traffic in a VLAN to a VSAN, as shown in Example 8-9.

Example 8-9 VLAN to VSAN Mapping

```
N5K-1# config
Enter configuration commands, one per line. End with CNTL/Z.
N5K-1(config)# vsan database
N5K-1(config-vsantdb)# vsan 2000 name FCOE
N5K-1(config-vsantdb)# vlan 2000
N5K-1(config-vlan)# fcoe vsan 2000
N5K-1(config-vlan)# end
```

```
N5K-1# show vlan fcoe
```

Original VLAN ID	Translated VSAN ID	Association State
-----	-----	-----
2000	2000	Operational

```
N5K-1#
```

Next, the Ethernet interface and VFC interface need to be configured to carry the Ethernet VLAN and VNP mode. Example 8-10 reflects this process.

Example 8-10 VNP Port Configuration on the Nexus 5000

```
N5K-1# config
Enter configuration commands, one per line. End with CNTL/Z.
N5K-1(config)# int e1/1
N5K-1(config-if)# switchport mode trunk
N5K-1(config-if)# switchport trunk allowed vlan 2000
N5K-1(config-if)# no shut
N5K-1(config-if)# desc FCoE-NPV Connection to N7K-1 E6/27
N5K-1(config-if)# interface vfc11
N5K-1(config-if)# desc FCoE-NPV Connection to N7K-1 vfc11
N5K-1(config-if)# switchport mode np
N5K-1(config-if)# bind interface e1/1
N5K-1(config-if)# switchport trunk allowed vsan 2000
N5K-1(config-if)# no shut
N5K-1(config-if)# end
N5K-1#N5K-1# show int vfc11
vfc11 is trunking
    Bound interface is Ethernet1/1
    Port description is FCoE-NPV Connection to N7K-1 vfc11
    Hardware is Ethernet
    Port WWN is 20:0a:00:05:73:d3:14:7f
    Admin port mode is NP, trunk mode is on
    snmp link state traps are enabled
    Port mode is TNP
    Port vsan is 1
    Trunk vsans (admin allowed and active) (2000)
    Trunk vsans (up) (2000)
    Trunk vsans (isolated) ()
    Trunk vsans (initializing) ()
    1 minute input rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
    1 minute output rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
    10 frames input, 1140 bytes
    0 discards, 0 errors
```



```

    7 frames output, 980 bytes
    0 discards, 0 errors
last clearing of "show interface" counters Mon Jul 30 17:21:52 2012

Interface last changed at Mon Jul 30 17:21:52 2012

```

```
N5K-1#
```

A similar configuration must be applied on the Nexus 7000 side of the link. The primary difference is that the VFC is configured for the VF mode and NPIV is enabled.

Example 8-11 shows the commands used for the configuration and the commands to verify the correct operation.

Example 8-11 VFC and Ethernet Port Configuration on the Nexus 7000

```

N7K-1-FCoE# config
Enter configuration commands, one per line. End with CNTL/Z.
N7K-1-FCoE(config)# feature npiv
N7K-1-FCoE(config)# interface Ethernet6/27
N7K-1-FCoE(config-if)# description FCoE-NPV Connection to N5K-1 e1/1
N7K-1-FCoE(config-if)# switchport
N7K-1-FCoE(config-if)# switchport mode trunk
N7K-1-FCoE(config-if)# switchport trunk allowed vlan 2000
N7K-1-FCoE(config-if)# no shutdown
N7K-1-FCoE(config-if)#
N7K-1-FCoE(config-if)#interface vfc11
N7K-1-FCoE(config-if)#bind interface Ethernet6/27
N7K-1-FCoE(config-if)# switchport trunk allowed vsan 2000
N7K-1-FCoE(config-if)# no shutdown
N7K-1-FCoE(config-if)#end
N7K-1-FCoE#N7K09-FCoE# show int vfc11
vfc11 is trunking
    Bound interface is Ethernet6/27
    Hardware is Ethernet
    Port WWN is 20:0a:00:26:98:0f:d9:bf
    Admin port mode is F, trunk mode is on
    snmp link state traps are enabled
    Port mode is TF
    Port vsan is 1
    Speed is auto
    Trunk vsans (admin allowed and active) (2000)
    Trunk vsans (up) (2000)
    Trunk vsans (isolated) ()
    Trunk vsans (initializing) ()

```

```

7 fcoe in packets
868 fcoe in octets
11 fcoe out packets
1324 fcoe out octets
Interface last changed at Mon Jul 30 17:44:30 2012

```

```
N7K01-FCoE# show fcns data
```

```
VSAN 2000:
```

```

-----
FCID          TYPE  PWWN                               (VENDOR)          FC4-TYPE:FEATURE
-----
0x010000      N     20:0a:00:05:73:d3:14:7f (Cisco)           npv

```

```
Total number of entries = 1
```

```
N7K-1-FCoE#
```

Nexus 7000 Unified Fabric Configuration

The Nexus 7000 provides director class support for FCoE solutions and can be used in both core and edge topologies. The platform provides the high-availability features and capabilities such as redundant supervisors, redundant hardware components, and the inherent availability components of NX-OS, such as Storage VDCs. In-Service Software Upgrade (ISSU), Stateful Switch Over (SSO) and stateful process restart make for a solid foundation.

FCoE on the Nexus 7000 is available on the F1 (N7K-F132XP-15) and F2/F2e (N7K-F248XP-25) modules. When using FCoE on the F2/F2e module, a Supervisor 2 or Supervisor 2E must be used. FCoE on F2/F2e cannot work with a Supervisor 1 module. FCoE is also a licensed feature, and the license is bound to a module, so if FCoE will be used across multiple modules in a chassis, there must be an FCoE license installed per module.

With these requirements met, FCoE can be installed on the Nexus 7000. FCoE installation requires the system QoS policy is configured to a template that provides a no-drop class. This is configured in either the default VDC or the admin VDC if running NX_OS 6.1(1) or later. The default QoS policy uses eight drop classes and is named **default-np-8e-policy**. Example 8-12 shows the QoS classes available to be selected and shows the change to a single no-drop class. This policy matches FCoE traffic in CoS 3 and provides a lossless Ethernet transport (no drop).

Example 8-12 *Setting the System QoS Policy*

```

N7K-1# config
Enter configuration commands, one per line. End with CNTL/Z.
N7K-1(config)# system qos

```

```

N7K-1(config-sys-qos)# service-policy type network-qos ?
  default-nq-4e-policy  Default 4-ethernet policy (4-drop 4-nodrop CoS)
  default-nq-6e-policy  Default 6-ethernet policy (6-drop 2-nodrop CoS)
  default-nq-7e-policy  Default 7-ethernet policy (7-drop 1-nodrop CoS)
  default-nq-8e-policy  Default 8-ethernet policy (8-drop CoS)
N7K-1(config-sys-qos)# service-policy type network-qos default-nq-7e-policy
N7K-1(config-sys-qos)# end
N7K-1# show policy-map system type network-qos

Type network-qos policy-maps
=====
policy-map type network-qos default-nq-7e-policy
  class type network-qos c-nq-7e-drop
    match cos 0-2,4-7
    congestion-control tail-drop
    mtu 1500
  class type network-qos c-nq-7e-ndrop-fcoe
    match cos 3
    match protocol fcoe
    pause
    mtu 2112

N7K-1#

```

With the QoS policy mapped to a no-drop policy, the next step is to install the FCoE feature set and configure a Storage VDC. This enables FCoE across the entire chassis and then creates a VDC to be used for storage functions. Example 8-13 describes this process.

Example 8-13 *Installing FCoE Feature Set and Creating a Storage VDC*

```

N7K-1# config
Enter configuration commands, one per line. End with CNTL/Z.
N7K-1(config)# install feature-set fcoe
N7K-1(config)# vdc FCoE type storage
Note: Creating VDC, one moment please ...
N7K-1(config-vdc)# show vdc

vdc_id  vdc_name  state  mac                type  lc
-----  -
1       N7K-1     active 00:26:98:0f:d9:c1  Admin  None
2       Agg1      active 00:26:98:0f:d9:c2  Ethernet f2
3       Core1     active 00:26:98:0f:d9:c3  Ethernet m1 f1 m1x1 m2x1
4       Access1   active 00:26:98:0f:d9:c4  Ethernet m1 f1 m1x1 m2x1
5       FCoE     active 00:26:98:0f:d9:c5  Storage  f1 f2

```

```

N7K-1(config-vdc)# show vdc FCoE detail

vdc id: 5
vdc name: FCoE
vdc state: active
vdc mac address: 00:26:98:0f:d9:c5
vdc ha policy: RESTART
vdc dual-sup ha policy: SWITCHOVER
vdc boot Order: 1
CPU Share: 5
CPU Share Percentage: 16%
vdc create time: Tue Jul 31 00:15:39 2012
vdc reload count: 0
vdc restart count: 0
vdc type: Storage
vdc supported linecards: f1 f2

N7K-1(config-vdc)#

```

The next step is to configure the storage VDC by allocating ports from modules, allocating a range of VLANs for use with FCoE, and then setting up the VDC for FCoE usage. Because this VDC is new, the switch prompts for a few items such as system password strength, password, and to run the setup script. When completed, basic FCoE configuration can begin. Example 8-14 walks through this process.

Example 8-14 *Allocation of Ports and Initial VDC Configuration*

```

N7K-1# config
Enter configuration commands, one per line. End with CNTL/Z.
N7K-1(config)# vdc fcoe
N7K-1(config-vdc)# allocate interface e6/17,e6/27,e6/29-32
Entire port-group is not present in the command. Missing ports will be included
automatically
Moving ports will cause all config associated to them in source vdc to be removed.
Are you sure you want to move the ports (y/n)? [yes] yes
N7K-1(config-vdc)# allocate fcoe-vlan-range 2000 from vdc Access1
N7K-1(config-vdc)# end
N7K-1#
N7K-1# switchto vdc fcoe

```

```
---- System Admin Account Setup ----
```

```
Do you want to enforce secure password standard (yes/no) [y]: n
```

```
Enter the password for "admin":
```

```
Confirm the password for "admin":
```

```
---- Basic System Configuration Dialog VDC: 5 ----
```

This setup utility will guide you through the basic configuration of the system. Setup configures only enough connectivity for management of the system.

Please register Cisco Nexus7000 Family devices promptly with your supplier. Failure to register may affect response times for initial service calls. Nexus7000 devices must be registered to receive entitled support services.

Press Enter at anytime to skip a dialog. Use ctrl-c at anytime to skip the remaining dialogs.

```
Would you like to enter the basic configuration dialog (yes/no): no
```

```
Cisco Nexus Operating System (NX-OS) Software
```

```
TAC support: http://www.cisco.com/tac
```

```
Copyright (c) 2002-2012, Cisco Systems, Inc. All rights reserved.
```

The copyrights to certain works contained in this software are owned by other third parties and used and distributed under license. Certain components of this software are licensed under the GNU General Public License (GPL) version 2.0 or the GNU Lesser General Public License (LGPL) Version 2.1. A copy of each such license is available at

```
http://www.opensource.org/licenses/gpl-2.0.php and
```

```
http://www.opensource.org/licenses/lgpl-2.1.php
```

```
N7K-1-FCoE#
```

```
N7K-1-FCoE# config
```

```
Enter configuration commands, one per line. End with CNTL/Z.
```

```
N7K-1-FCoE(config)# feature-set fcoe
```

```
N7K-1-FCoE(config)# feature npiv
```

```
N7K-1-FCoE(config)# feature lldp
```

```
N7K-1-FCoE(config)# vsan database
```

```
N7K-1-FCoE(config-vsantdb)# vsan 2000
```

```
N7K-1-FCoE(config-vsantdb)# vlan 2000
```

```

N7K-1-FCoE(config-vlan)# fcoe
N7K-1-FCoE(config-vlan)# end
N7K-1-FCoE(config)# end
N7K-1-FCoE#
N7K-1-FCoE# show vlan fcoe

Original VLAN ID          Translated VSAN ID      Association State
-----
                2000                2000                Operational
N7K-1-FCoE#
N7K-1-FCoE(config)# end
N7K-1-FCoE#

```

With the foundation for FCoE configured, the next step is to provision connectivity. Figure 8-13 shows the topology the following examples use.




Figure 8-13 FCoE Topology Between Nexus 7000 and MDS

The first step is to configure the Ethernet interfaces, add them to a port channel for additional bandwidth on the ISL and redundancy, and then configure the VFC, as shown in Example 8-15.

Example 8-15 *Nexus 7000 to MDS Interconnection*

```

N7K-1-FCoE# config
Enter configuration commands, one per line. End with CNTL/Z.
N7K-1-FCoE(config)# feature lacp
N7K-1-FCoE(config)# int e6/29-32
N7K-1-FCoE(config-if-range)# channel-group 258 mode active
N7K-1-FCoE(config-if-range)# int po258
N7K-1-FCoE(config-if)# desc Port Channel to MDS9506-1
N7K-1-FCoE(config-if)# switchport mode trunk
N7K-1-FCoE(config-if)# switchport trunk allowed vlan 2000
N7K-1-FCoE(config-if)# no shut
N7K-1-FCoE(config-if)# int vfc 101
N7K-1-FCoE(config-if)# switchport desc VE Port Channel to MDS9506-1
N7K-1-FCoE(config-if)# switch mode e
N7K-1-FCoE(config-if)# switch trunk allowed vsan 2000
N7K-1-FCoE(config-if)# bind interface po258
N7K-1-FCoE(config-if)# no shut
N7K-1-FCoE(config-if)# end
N7K-1-FCoE# show int vfc101
vfc101 is trunking
    Bound interface is port-channel258
    Port description is VE Port Channel to MDS9506-1
    Hardware is Ethernet
    Port WWN is 20:64:00:26:98:0f:d9:bf
    Admin port mode is E, trunk mode is on
    snmp link state traps are enabled
    Port mode is TE
    Port vsan is 1
    Speed is 40 Gbps
    Trunk vsans (admin allowed and active) (2000)
    Trunk vsans (up) (2000)
    Trunk vsans (isolated) ()
    Trunk vsans (initializing) ()
    120677 fcoe in packets
    13910628 fcoe in octets
    120679 fcoe out packets
    10352660 fcoe out octets
    Interface last changed at Tue Jul 31 01:21:17 2012

N7K-1-FCoE#

```

For reference, Example 8-16 shows the corresponding configuration on the MDS.

Example 8-16 *MDS FCoE Configuration*

```

MDS9506-1# show run int Eth3/5, Eth3/6, Eth4/5, Eth4/6

!Command: show running-config interface Ethernet3/5-6, Ethernet4/5-6
!Time: Tue Jul 31 01:30:57 2012

version 5.2(2a)

interface Ethernet3/5
  switchport mode trunk
  switchport trunk allowed vlan 2000
  channel-group 258 mode active
  no shutdown

interface Ethernet3/6
  switchport mode trunk
  switchport trunk allowed vlan 2000
  channel-group 258 mode active
  no shutdown

interface Ethernet4/5
  switchport mode trunk
  switchport trunk allowed vlan 2000
  channel-group 258 mode active
  no shutdown

interface Ethernet4/6
  switchport mode trunk
  switchport trunk allowed vlan 2000
  channel-group 258 mode active
  no shutdown

MDS9506-1#
MDS9506-1# show run int epo258

!Command: show running-config interface ethernet-port-channel258
!Time: Tue Jul 31 01:31:42 2012

version 5.2(2a)

interface ethernet-port-channel258
  switchport mode trunk
  switchport trunk allowed vlan 2000

```



```

Invalid interface format at '^' marker.
MDS9506-1# show run int vfc101

!Command: show running-config interface vfc101
!Time: Tue Jul 31 01:31:52 2012

version 5.2(2a)

interface vfc101
  bind interface ethernet-port-channel258
  switchport mode E
  switchport trunk allowed vsan 2000
  no shutdown

MDS9506-1# MDS9506-1# show int vfc101
vfc101 is trunking
  Bound interface is ethernet-port-channel258
  Hardware is Ethernet
  Port WWN is 20:64:00:0d:ec:35:1e:ff
  Admin port mode is E, trunk mode is on
  snmp link state traps are enabled
  Port mode is TE
  Port vsan is 1
  Speed is 40 Gbps
  Trunk vsans (admin allowed and active) (2000)
  Trunk vsans (up) (2000)
  Trunk vsans (isolated) ()
  Trunk vsans (initializing) ()
  117696 fcoe in packets
  10091312 fcoe in octets
  117695 fcoe out packets
  13575440 fcoe out octets
  Interface last changed at Tue Jul 31 01:17:09 2012

MDS9506-1#

```

FCoE on the Nexus 7000 also supports a unique capability that enables interfaces to be shared between two VDCs. This enables the Nexus 7000 to be used in the access layer of networks where servers connect to the switch and use FCoE. A shared interface enables FCoE traffic to be segmented into the Storage VDC at the edge of the network. When an interface is shared between two VDCs, a few rules must be followed:

- Interfaces can be shared only between one Ethernet VDC and one Storage VDC.
- Interfaces to be shared must be configured as 802.1Q trunks in the Ethernet VDC.
- Interfaces may be shared only from the Ethernet VDC that allocated VLANs to the Storage VDC.
- The Ethernet VDC “owns” the physical interface. If the interface is admin down in the Ethernet VDC, it will be admin down in the Storage VDC.
- All ports that have a common ASIC must be allocated as shared interfaces. This is done in groups of two on the F1 modules and groups of four on F2/F2e modules.

Note If all the shared ASIC ports are not configured as trunks, allocation as shared interfaces will fail.

In Example 8-17, four ports are configured as trunks in the Ethernet VDC and then configured for shared interfaces in the Storage VDC.

Example 8-17 *Nexus 7000 Shared Interface Allocation*

```
N7K-1# config
Enter configuration commands, one per line. End with CNTL/Z.
N7K-1(config)# vdc fcoe
N7K-1(config-vdc)# allocate shared interface e6/17
Entire port-group is not present in the command. Missing ports will be included
automatically

Ports that share the port group of the interfaces you have specified will be affected
as well. Continue (y/n)? [yes] yes
N7K-1(config-vdc)# end
N7K-1# fcoe
Cisco Nexus Operating System (NX-OS) Software
TAC support: http://www.cisco.com/tac
Copyright (c) 2002-2012, Cisco Systems, Inc. All rights reserved.
The copyrights to certain works contained in this software are
owned by other third parties and used and distributed under
license. Certain components of this software are licensed under
the GNU General Public License (GPL) version 2.0 or the GNU
Lesser General Public License (LGPL) Version 2.1. A copy of each
such license is available at
http://www.opensource.org/licenses/gpl-2.0.php and
http://www.opensource.org/licenses/lgpl-2.1.php
N7K-1-FCoE# show int brief
```

```
-----
Interface                Status                Speed
```

```

-----
sup-fc0                               (Gbps)
up                                     1
-----
Ethernet      VLAN      Type Mode      Status Reason                               Speed      Por
t
Interface
#
-----
Eth6/17       1          eth trunk  down    Administratively down    auto(D)  --
Eth6/18       1          eth trunk  down    Administratively down    auto(D)  --
Eth6/19       1          eth trunk  down    Administratively down    auto(D)  --
Eth6/20       1          eth trunk  down    Administratively down    auto(D)  --
Eth6/25       --         eth routed down    Administratively down    auto(D)  --
Eth6/26       --         eth routed down    Administratively down    auto(D)  --
Eth6/27       --         eth routed down    Administratively down    auto(D)  --
Eth6/28       --         eth routed down    SFP not inserted        auto(D)  --
Eth6/29       1          eth trunk  up      none                      10G(D)  258
Eth6/30       1          eth trunk  up      none                      10G(D)  258
N7K-1-FCoE#

```

The next step required is to create the VFC interface for the host and specify the shared interface as the binding. This is the same syntax used on the Nexus 5x00 earlier in the chapter. Example 8-18 shows the process for the topology shown in Figure 8-14.




Figure 8-14 FCoE Topology

Example 8-18 *Nexus 7000 VFC Interface Creation*

```

N7K-1-FCoE# config
Enter configuration commands, one per line. End with CNTL/Z.
N7K-1-FCoE(config)# int vfc617
N7K-1-FCoE(config-if)# bind interface ethernet 6/17
N7K-1-FCoE(config-if)# no shut
N7K-1-FCoE(config-if)# end
N7K-1-FCoE#

```

With the VFCs created and bound, VEs created to the MDS, and both storage and hosts connected to the fabric, the last step would be to configure zoning and device aliasing for the FC network. The Nexus switches can participate in zoning with a Fibre Channel network.

Note Nexus 5x00 when configured in the NPV or FCoE-NPV mode do not participate in zoning and aliasing because they rely on the upstream device to perform those functions.

Example 8-19 shows a device-alias, zone and zoneset creation, and activation.

Example 8-19 *Device alias, zone, and zoneset Creation and Activation*

```

N7K-1-FCoE# config
Enter configuration commands, one per line. End with CNTL/Z.
N7K-1-FCoE(config)# device-alias mode enhanced
N7K-1-FCoE(config)# device-alias database
N7K-1-FCoE(config-device-alias-db)# device-alias name C210-ESX1 pwwn
20:00:e8:b7:48:4d:74:22
N7K-1-FCoE(config-device-alias-db)# device-alias name NetApp_FAS270 pwwn
50:0a:09:81:85:75:90:88
N7K-1-FCoE(config-device-alias-db)# device-alias commit
N7K-1-FCoE(config-device-alias-db)# exit
N7K-1-FCoE(config)#zone name NetappArray vsan 2000
N7K-1-FCoE(config-zone)# member device-alias NetApp_FAS270
N7K-1-FCoE(config-zone)# member device-alias C210-ESX1
N7K-1-FCoE(config-zone)# zone name C210-ESX1 vsan 2000
N7K-1-FCoE(config-zone)# member device-alias C210-ESX1
N7K-1-FCoE(config-zone)# zoneset name VSAN2000_ZS vsan 2000
N7K-1-FCoE(config-zoneset)# member NetappArray
N7K-1-FCoE(config-zoneset)# member C210-ESX1
N7K-1-FCoE(config-zoneset)# exit
N7K-1-FCoE(config)# zoneset activate name VSAN2000_ZS vsan 2000
N7K-1-FCoE(config)# end
N7K-1-FCoE#

```

Summary

Unified Fabric offers several benefits to customers, including

- **Lower capital expenditures:** Through the reduction of adapters, cables, and ports required within the infrastructure.
- **Lower operational expenses:** Through the reduction of adapters, cables, and ports drawing power within the data center.
- **Reduced deployment cycles:** Provides a wire-once model, where all LAN, SAN, IPC, and management traffic is available to every server without requiring additional connectivity components.
- **Higher availability:** Few adapters and fewer ports means fewer components that could fail.

By taking advantage of enhancements to traditional Ethernet technologies, and the emergence of technologies such as FCoE, customers can realize these benefits with minimal disruption to operational models. This chapter showed the basic Nexus 5x00 and Nexus 7000 configurations necessary to provide a Unified access method for LAN data traffic and SAN storage traffic. The multiple technologies that can be used with Unified Fabric such as NPV, NPIV FCoE-NPV, Storage VDCs, and shared interfaces were illustrated, and various use cases were discussed.

Index

NUMBERS

6PE and MPLS (Multiprotocol Label Switching), 725

6VPE and MPLS (Multiprotocol Label Switching), 725

10-Gigabit Ethernet, 456-458

A

AAA (Authentication, Authorization, and Accounting)

accounting, 256-257

authentication, 256, 258-259, 263, 265

authorization, 256

RADIUS, 258-259

configuration distribution, 259-264

configuring, 256-259

TrustSec

AAA configuration, 281-294

authentication methods, 279

authorization methods, 279

access switches, PVLAN (Private Virtual Local Area Networks) configuration, 78-79

accounting, 256-257

ACL (Access Control Lists), 302-305

MAC ACL, 305-307

QoS, 645

VLAN ACL, 307-308

AED (Authoritative Edge Devices) and OTV (Overlay Transport Virtualization), 680-681

alert groups (Smart Call Home), 424

Alphanumeric Process ID, OSPF

network advertisements, 157-158

passive interfaces, 159

stub routing, 164-166

summarization, 161

alphanumeric strings

EIGRP instances, creating, 138

OSPF, 156

alphanumeric tags

EIGRP

passive interfaces, 142

stub routing, 146

summarization, 143

- SVI, adding to EIGRP, 140-141
- analyzer (embedded), 412-414**
- Anycast-RP and RP configuration, 227, 237**
- ARP (Address Resolution Protocol)**
 - ARP resolution verification in FabricPath, 125-126
 - Dynamic ARP
 - configuring, 316-317*
 - Inspection Trust state, 317-320*
 - OTV, 700-702
 - synchronization, port channels and, 117
- ASM (Any Source Multicast), 225**
- Atomic configuration rollbacks, 33**
- Atomic mode (checkpoints)**
 - Nexus 5x00, 434
 - Nexus 7000, 431
- authentication, 256**
 - BGP, 190-192
 - EIGRP, 147-149
 - GLBP, 215-217
 - HSRP, 174-203
 - OSPF, 167-169
 - PIM, 229-230
 - RADIUS, 258-259, 263, 265
 - SNMPv3, 340
 - TrustSec, 279
 - VRRP, 208-209
- authorization, 256, 279**
- autonomous system numbers, EIGRP configuration, 138-139**
- Auto-RP and RP configuration, 227, 235-237**
- availability (high)**
 - GOLD, 358-359
 - bootup diagnostics, 359*

- on-demand diagnostics, 365*
- runtime diagnostics, 360-364*
- ISSU, 370-383
- nonstop forwarding, 370
- NX-OS, 366
- processes
 - modularity, 366-368*
 - restarts, 368-369*
- redundancy (physical), 349-350
 - cooling systems, 352-355*
 - EOBC, 357*
 - fabric modules, 357-358*
 - power supplies, 350-352*
 - supervisors, 355-357, 369-370*
- stateful switchovers, 368-369
- Unified Fabric, 456, 488

B

- B22HP HP-FEX, NX-OS support, 7**
- Best-effort configuration rollbacks, 33**
- Best-Effort mode mode (checkpoints)**
 - Nexus 5x00, 434
 - Nexus 7000, 431
- BGP (Border Gateway Protocol)**
 - authentication, 190-192
 - configuring, 183-185
 - dotted decimal notation, 185-186*
 - plain text notation, 185*
 - Router ID, 186*
 - routing families, 186-187*
 - verifying configurations, 198*
 - enabling, 184-185
 - neighbors, 187-190

- network advertisements, 194-196
- overview of, 183
- peer templates, 192-194
- routing metrics, modifying, 197
- security, 190-192
- topology of, 184
- BiDir (Bidirectional shared trees),
226, 234-235, 237
- Bootflash files, syntax of, 30
- bootup diagnostics, 359
- BPDU (Bridge Protocol Data Units),
OTV failure isolation, 698-699
- BPDUFilter, 95-96
- BPDUGuard, 94-95
- Bridge Assurance, 98-99
- broadcast storms, traffic storm
control, 323-324
- broadcast suppression, PVLAN, 76
- BSR (Boot Strap Routers), RP
configuration, 226-227, 232-235

C

- case studies, Nexus migration
 - design, 751
 - access layer, 751*
 - aggregation layer, 751*
 - Core layer, 751*
 - design goals, 750
 - availability, 751*
 - flexibility, 751*
 - performance, 750-751*
 - scalability, 750*
 - existing data center environment,
749-750
 - goal of, 749
 - on-going maintenance windows, 788

- maintenance window #1, 754
 - port channel builds, 754-757*
 - port channel configuration,
757-758*
 - port channel verification,
758-759*
 - summary, 760*
- maintenance window #2, 760-761
 - bandwidth configuration, 779*
 - ECMP verification, 776-777*
 - EIGRP, 776*
 - enabling port channels, 780*
 - FEX configuration, 786-787*
 - HSRP routers, 777-779*
 - port channel configuration,
761-763*
 - port channel shut downs, 780*
 - port channel verification,
765-766*
 - Rapid-PVST+766*
 - route verification, 780*
 - routed uplink shut downs, 782*
 - routing table verification,
775-776*
 - STP configuration, 763-764,
766*
 - STP verification, 766-768,
780-781*
 - SVI address configuration, 783*
 - SVI configuration, 769-774*
 - SVI shut downs, 781-782*
 - SVI verification, 775*
 - VLAN configuration, 763-764,
766*
 - vPC configuration, 764-765,
768, 784-786*
- migration plan, 752
- premigration steps, 752-754

- summary of, 788
- topology of, 760, 768, 782-784
- CE (Customer Edge) routers and MPLS (Multiprotocol Label Switching), 719
- checkpoints
 - Nexus 5x00, 434-435
 - Atomic mode*, 434
 - Best-Effort mode*, 434
 - creating*, 435-437
 - rollbacks*, 437
 - Stop-at-First-Failure mode*, 434
 - Nexus 7000, 431-432
 - Atomic mode*, 431
 - Best-Effort mode*, 431
 - creating*, 432-433
 - rollbacks*, 434
 - Stop-at-First-Failure mode*, 431
- CMP (Connectivity Management Processor) management interface, 15-18
- community VLAN (Virtual Local Area Networks), 76-77
- configuration files
 - Atomic configuration rollbacks, 33
 - Best-effort configuration rollbacks, 33
 - rollbacks
 - creating*, 34-35
 - limitations of*, 34
 - Stop-at-first-failure configuration rollbacks, 33
 - Verbose mode configuration rollbacks, 33
- control plane
 - MPLS, 718
 - OTV, 682-687

- Control VLAN (Virtual Local Area Networks), Nexus 1000V and, 492-493
- cooling systems, redundancy, 352-355
- CoPP (Control Plane Policing)
 - configuring, 327-335
 - multicasts, 253
- CP (Control Processor)/Supervisor management interface, 15-16
- cut-through switching, 60

D

- data packet rate limits, 335-340
- DCNM (Data Center Network Manager), 26-28
- debug commands, troubleshooting via, 54-56
- debug files, syntax of, 30
- DHCP (Dynamic Host Configuration Protocol) snooping, 313-316
- diagnostics, GOLD, 358-359
 - bootup diagnostics, 359
 - on-demand diagnostics, 365
 - runtime diagnostics, 360-364
- dispute mechanism, 98
- dotted decimal notation, BGP configuration, 185-186
- DVS (Distributed Virtual Switches). *See* Nexus 1000V
- Dynamic ARP (Address Resolution Protocol)
 - configuring, 316-317
 - Inspection Trust state, 317-320

E

edge ports (Spanning Tree Protocol), 99

EEE (Energy Efficient Ethernet), 447-448

EID (End-point Identifier) addresses, 730

EIGRP (Enhanced Interior Gateway Routing Protocol)

authentication, 147-149

configuring, 137-142

autonomous system numbers, 138-139

IPv4/IPv6 address families, 139

passive interfaces, 141-142

wide metrics, 141

enabling, 137

instances, creating via

alphanumeric strings, 138

numeric tags, 137-138

operation of, 136

overview of, 135-136

redistribution, 149

default routes, 149-151

Fixed Limit option, 152

prefix limits, 152

prefix list definitions, 150

reviewing, 153-154

route map definitions, 150

routing detail, 152-153

Warning option, 152

Withdraw option, 152-153

security, 147-149

stub routing, 145-147

summarization, 142-145

SVI, adding via

alphanumeric tags, 140-141

numeric tags, 139-140

VRF configuration, 714-715

embedded analyzer, 414-423

embedded serviceability

checkpoints

Nexus 5x00, 431-434

Nexus 7000, 431-434

EEE, 447-448

embedded analyzer, 412-414

ERSPAN

Nexus 5x00, 412-414

Nexus 1000V, 400-406

Nexus 7000, 406-412

NetFlow, 437-438

Nexus 1000V configuration, 442-444

Nexus 7000 configuration, 438-442

NTP, 444-445

POAP, 448

PTP, 445-447

Python, 449-454

Smart Call Home, 424-428

alert groups, 424

configuring, 424-431

executed commands, 424

SPAN

Nexus 5x00, 392-393

Nexus 1000V, 397-398

Nexus 7000, 386-392

overview of, 386

encryption, SNMPv3, 340

EOBC (Ethernet Out-of-Band), redundancy, 357

errdisable action (port security violations), 311

errdisable recovery, 95

ERSPAN (Encapsulated Remote SPAN)

Nexus 5x00, 412-414

Nexus 1000V, 400-406

Nexus 7000, 406-412

ESM (Extended Subnet Mode), LISP host mobility, 736-745

ESX vSphere hosts, adding VEM to, 528-535

Etherchannel

Layer 2 configurations, 61

Nexus 2000, 66-69

Ethernet, 456

10-Gigabit Ethernet, 456-458

EEE, 447-448

FCoE, 41-42, 458-460

Multihop FCoE, 462-463

Nexus 5x00 United Fabric configuration, 467-468

FCoE NPV mode, 466

lvEth ports, 495

Single-Hop FCoE, 461-462

Nexus 5x00 United Fabric configuration, 469-473

VEM in Nexus 1000V, 490, 492-494

functions of, 490-491

port profiles, 542-552

VEM additions to ESX vSphere hosts, 528-535

VEM installation via Installation Management Center, 519-525

VEM physical port classifications, 495

VEM supported ports, 494

VEM virtual port classifications, 494-495

vEth ports, 494

ETR (Egress Tunnel Routers), 730, 734-735

PETR, 731

xTR, 730

EXEC Command Mode, 12-13

executed commands (Smart Call Home), 424

F

Fabric Extenders (Nexus 2000), 60-61

fabric modules, redundancy, 357-358

FabricPath, 2

adjacencies, verifying, 124-125

ARP resolution, verifying, 125-126
configuring

interfaces, 124

switch IDs, 122

VLAN, 123-124

connectivity, testing, 125

MAC address tables, verifying, 126

Nexus 7000, enabling on, 122

overview of, 119-122

routing tables, displaying, 126-127

topology of, 122-123

vPC+, 127-132

FC uplinks, Nexus 5x00 United Fabric configuration, 470-471

FCoE (FibreChannel over Ethernet), 41-42, 458-460

Multihop FCoE, 462-463

- Nexus 5x00 United Fabric configuration
 - enabling FCoE*, 468
 - FCoE NPV mode*, 466, 473-477
 - Single-Hop FCoE*, 461-462, 469-473
 - verifying FCoE licenses*, 467-468
 - Nexus 7000 United Fabric configuration, 478-479
 - Nexus 7000/MDS topology, 481-484
 - FEX (Fabric Extenders)**
 - Nexus 5x00 QoS, 668-669
 - Nexus 7000 QoS, 661-662
 - FHRP (First Hop Redundancy Protocols)**, 198
 - GLBP, 198
 - authentication*, 215-217
 - configuring*, 212-214
 - enabling*, 212
 - overview of*, 211-212
 - priority and preempt*, 213-214
 - secondary IP addresses/subnets*, 217-219
 - security*, 215-217
 - HSRP, 198
 - authentication*, 174-203
 - configuring*, 199-201
 - IPv6*, 204-205
 - overview of*, 199
 - priority and preempt*, 200-201
 - secondary IP addresses/subnets*, 203-204
 - security*, 174-203
 - vPC interactions*, 210-211
 - OTV FHRP localization, 702-704
 - VRRP, 198
 - authentication*, 208-209
 - configuring*, 205-208
 - enabling*, 206
 - overview of*, 205
 - priority and preempt*, 207
 - secondary IP addresses/subnets*, 209
 - security*, 208-209
 - vPC interactions*, 210-211
 - files**
 - Bootflash files, 29
 - debug files, 29
 - Log files, 29
 - Nvram files, 29
 - slot0 files, 29
 - system files, 29
 - system files management, 28
 - configuration files*, 33-35
 - file systems*, 28-33
 - operating system files*, 35-36
 - usb1 files, 29
 - usb2 files, 30
 - volatile files, 29
 - Fixed Limit option**
 - EIGRP redistribution, 152
 - OSPF redistribution, 174
 - forwarding (nonstop)**, 370
 - Forwarding table (MPLS)**, 718
 - further reading, NX-OS**, 57
- ## G
-
- GLBP (Global Load Balancing Protocol)**, 198
 - authentication, 215-217
 - configuring, 212-214
 - enabling, 212

- OTV GLBP localization, 702-704
- overview of, 211-212
- priority and preempt, 213-214
- secondary IP addresses/subnets, 217-219
- security, 215-217
- topology of, 212

Global Configuration Command Mode, 13

GOLD (Generic Online Diagnostics), 358-359

- bootstrap diagnostics, 359
- on-demand diagnostics, 365
- runtime diagnostics, 360-364

H

HA policies, 54

Hello authentication, PIM (Protocol Independent Multicast), 229-230

help

- debug commands, troubleshooting via, 54-56
- MPLS, 725
- show commands, troubleshooting via, 54-55

high availability

- GOLD, 358-359
 - bootstrap diagnostics, 359*
 - on-demand diagnostics, 365*
 - runtime diagnostics, 360-364*

- ISSU, 370-383

- nonstop forwarding, 370

- NX-OS, 366

- processes

- modularity, 366-368*
 - restarts, 368-369*

- redundancy (physical), 349-350
 - cooling systems, 352-355*
 - EOBC, 357*
 - fabric modules, 357-358*
 - power supplies, 350-352*
 - supervisors, 355-357, 369-370*

- stateful switchovers, 368-369

- Unified Fabric, 456, 488

hostnames, changing in VSM, 536

HSRP (Hot-Standby Routing Protocol), 198

- authentication, 174-203
- configuring, 199-201
- enabling, 199
- IPv6, 204-205
- OTV HSRP localization, 702-704
- overview of, 199
- priority and preempt, 200-201
- secondary IP addresses/subnets, 203-204
- security, 202-203
- vPC interactions, 210-211

hypervisor, Nexus 1000V and, 489-490

IGMP (Internet Group Management Protocol)

configuring

- Nexus 1000V configuration, 246-248*

- Nexus 5000 configuration, 245*

- Nexus 7000 configuration, 242-245*

- enabling, 243

- overview of, 241-242

- snooping
 - Nexus 1000V snooping*, 246
 - Nexus 7000/Nexus 5000 snooping*, 244-245
 - vEthernet1 snooping*, 247-248
- topologies
 - Nexus 1000V configuration*, 247
 - Nexus 7000 configuration*, 242
- Inspection Trust state (Dynamic ARP)**, 317-320
- Interface Configuration Command Modd**, 13-14
- IOS**
 - NX-OS comparison, 10-12
 - OSPF, 158
- IP addresses and PVLAN (Private Virtual Local Area Networks)**, 76
- IP multicast**
 - CoPP, 253
 - IGMP, 241-242
 - enabling*, 243
 - Nexus 1000V configuration*, 246-248
 - Nexus 1000V snooping*, 246
 - Nexus 5000 configuration*, 245
 - Nexus 7000 configuration*, 242-245
 - Nexus 7000/Nexus 5000 snooping*, 244-245
 - vEthernet1 snooping*, 246
- multicast distribution trees, 222-223
- overview of, 221-222
- PIM, 225-226
 - ASM, 225
 - BiDir*, 226, 234-235, 237
 - enabling*, 228
 - join policy configurations*, 252-253
 - MSDP, 226, 248-250
 - Nexus 7000/Nexus 5000 configuration*, 227-241
 - prune policy configurations*, 252-253
 - RP administrative scoping*, 250-251
 - SSM, 225, 239-241
 - static RPF modes*, 226, 239-241
 - topology of*, 227, 237
- reverse path forwarding, 225
- RP, 226-227
- IP Source Guard**, 321-322
- IPv4**
 - address families, configuring for EIGRP, 139
 - BGP neighbors, 187-190
 - IS-IS configuration, 180
- IPv6**
 - address families, configuring for EIGRP, 139
 - IS-IS configuration, 181
 - MPLS, 725
- IS-IS**
 - configuring, 178-183
 - IPv4-enabled instances*, 180
 - IPv6-enabled instances*, 181
 - enabling, 179
 - MPLS TE configuration, 724
 - overview of, 178
 - VRF configuration, 715
- isolated VLAN (Virtual Local Area Networks)**, 76
- ISSU (In-Service Software Upgrades)**, 370-383

ITR (Ingress Tunnel Routers), 730,
734-735
PITR, 730
xTR, 730

J - K - L

join policies in PIM configurations,
252-253

keychain management, 322-323

labels, defining in MPLS
(Multiprotocol Label Switching),
718

LACP (Link Aggregation Control
Protocol), 104-107

Layer 2

cut-through switching, 60

enabling, 77

FabricPath

adjacency verification, 124-125

ARP resolution verification,
125-126

connectivity tests, 125

enabling on Nexus 7000, 122

interface configuration, 124

*MAC address table
verification*, 126

overview of, 119-122

routing table verification,
126-127

switch ID configuration, 122

topology of, 122-123

VLAN configuration, 123-124

vPC+127-132

MAC addresses

checking table consistency, 71

clearing tables, 70

displaying tables, 70

Nexus 2000

Etherchannel, 61, 66-69

Fabric Extenders, 60-61

port-channel forwarding, 61,
66-69

static pinning, 61-66

switching, 61

overview of, 59-60

port channels

ARP synchronization, 117

flow control, 107-108

*LACP physical port
assignments*, 104-107

load balancing, 103-104

*Nexus 7000 port channel
forwarding*, 69-71

overview of, 103

*static physical port
assignments*, 104-105

verifying load distribution,
108-109

vPC, 109-117

port profiles, 102-103

PVLAN

broadcast suppression, 76

community VLAN, 76-77

configuring access switches,
78-79

creating SVI, 78

defining, 78

enabling, 77

IP addresses, 76

isolated VLAN, 76

overview of, 76-77

primary VLAN, 76

secondary VLAN, 76-77

security, 76

- uses for*, 76
- verifying configurations*, 80
- standard hosts, configuring, 101
- store-and-forward switching, 60
- STP, 80
 - BPDUFilter*, 95-96
 - BPDUGuard*, 94-95
 - Bridge Assurance*, 98-99
 - defining default port type*, 99-100
 - dispute mechanism*, 98
 - edge ports*, 99
 - errdisable recovery*, 95
 - LoopGuard*, 97-98
 - MST, 81, 87-91
 - network ports*, 99
 - normal ports*, 99
 - port costs*, 91-94
 - port priority*, 94
 - Rapid-PVST*, 80-87
 - RootGuard*, 96-97
 - timers*, 84-87
 - TrunkFast*, 100
 - virtualization hosts*, 100-103
 - VLAN, 81
- trunk ports, configuring, 100-101
- UDLD, 118-119
- virtualization hosts, linking to, 101
- VLAN
 - adding ports*, 73
 - adding/removing from trunks*, 74
 - assigning memberships*, 73-74
 - configuring*, 72
 - configuring trunk interfaces*, 74
 - creating*, 72
 - creating multiple VLAN*, 72

- displaying internal VLAN*, 71-72
- overview of*, 71
- Spanning Tree Protocol*, 81
- verifying configurations*, 74-76
- VTP, 72-73

Layer 3, 179

BGP

- authentication*, 190-192
- configuring*, 183-187
- enabling*, 184-185
- neighbors*, 187-190
- network advertisements*, 194-196
- overview of*, 183
- peer templates*, 192-194
- routing metric modifications*, 197
- security*, 190-192
- verifying configurations*, 198

EIGRP

- authentication*, 147-149
- configuring*, 137-142
- enabling*, 137
- instance creation*, 137-138
- operation of*, 136
- overview of*, 135-136
- redistribution*, 149-154
- security*, 147-149
- stub routing*, 145-147
- summarization*, 142-145
- SVI additions*, 139-141

GLBP, 198

- authentication*, 215-217
- configuring*, 212-214
- enabling*, 212
- overview of*, 211-212

- priority and preempt*, 213-214
- secondary IP addresses/
subnets*, 217-219
- security*, 215-217
- HSRP, 198
 - authentication*, 174-203
 - configuring*, 199-201
 - IPv6, 204-205
 - overview of*, 199
 - priority and preempt*, 200-201
 - secondary IP addresses/
subnets*, 203-204
 - security*, 202-203
 - vPC interactions*, 210-211
- IS-IS
 - configuring*, 178-183
 - enabling*, 179
 - overview of*, 178
- MPLS
 - 6PE, 725
 - 6VPE, 725
 - CE, 719
 - control plane*, 718
 - device relationships*, 719
 - enabling*, 721
 - Forwarding table*, 718
 - high availability*, 725-726
 - IPv6, 725
 - labels*, 718
 - Layer 3 VPN, 718, 720-723
 - LDP, 719, 721-723
 - LSP, 719
 - LSR, 719
 - managing*, 725
 - Nexus hardware requirements*,
726-727
 - NX-OS licensing*, 726-727
- overview of*, 717-718
- P, 719
- PE, 719
- Quality of Service*, 718, 723
- RD, 719
- RSVP, 719
- RT, 719
- Traffic Engineering*, 718,
723-725
- troubleshooting*, 725
- VRF-based CE-PE BGP
configuration*, 723
- Nexus 1000V VSM installation,
536-542
- OSPF
 - authentication*, 167-169
 - configuring*, 154-160
 - OSPFv3 configuration*,
177-178
 - overview of*, 154
 - redistribution*, 169-177
 - security*, 167-169
 - stub routing*, 163-167
 - summarization*, 160-163
- VRF
 - EIGRP configuration*, 714-715
 - Interface and VRF
configuration*, 712
 - IS-IS configuration*, 715
 - MPLS, VRF-based CE-PE BGP
configuration*, 723
 - Nexus hardware requirements*,
726-727
 - NX-OS licensing*, 726-727
 - operational commands*, 713
 - OSPF configuration*, 714
 - overview of*, 709
 - predefined VRF*, 710-713

- routing table isolation, 712-713*
 - VRF-capable components, 710*
 - VRF-Lite, 713-717*
 - VRRP, 198
 - authentication, 208-209*
 - configuring, 205-208*
 - enabling, 206*
 - overview of, 205*
 - priority and preempt, 207*
 - secondary IP addresses/subnets, 209*
 - security, 208-209*
 - vPC interactions, 210-211*
 - LDP (Label Distribution Protocol), MPLS, OSPF configuration, 722
 - LDP (Label Distribution Protocol) and MPLS (Multiprotocol Label Switching), 719, 722
 - enabling LDP, 721-722
 - OSPF configuration, 721-722
 - licensing NX-OS, 7-10
 - LinkSec, TrustSec configurations, 294-302
 - LISP (Locator/ID Separator Protocol), 730, 746-747
 - communication sequence of events, 733
 - components of, 731
 - control plane sequence of events, 732
 - data plane, 733-735
 - deploying, 746
 - EID addresses, 730
 - ESM, 736-745
 - ETR, 730, 734-735
 - hardware requirements, 731-732
 - host mobility, 736-745
 - ITR, 730, 734-735
 - LISP/non-LISP site communications, 735-736
 - MR, 730
 - MS, 730
 - OTV, 707
 - overview of, 729-730
 - PETR, 731
 - PITR, 730
 - RLOC, 736, 738
 - RLOC addresses, 730
 - software requirements, 731-732
 - xTR, 730
 - Log files, syntax of, 30
 - loop prevention, Nexus 1000V and, 495
 - LoopGuard, 97-98
 - Loose Unicast RPF mode (unicast RPF), 326
 - LSP (Label Switching Path) and MPLS (Multiprotocol Label Switching), 719
 - LSR (Label Switching Routers) and MPLS (Multiprotocol Label Switching), 719
 - lvEth (local virtual Ethernet) ports, VEM in Nexus 1000V, 495
- ## M
-
- MAC ACL (Access Control Lists), 305-307
 - MAC addresses
 - FabricPath, verifying MAC address tables, 126
 - Layer 2
 - checking table consistency, 71*
 - clearing tables, 70*
 - displaying tables, 70*
 - OTV Edge Device, 689

machine managers (virtual). *See*
hypervisor

management interfaces, 14-28

MDS 9000, NX-OS support, 7

MDS ports, Nexus 5x00 United
Fabric configuration, 470

memberships (VLAN), assigning,
73-74

MGMT0 management interface, 15

migration case study (Nexus)

design, 751

access layer, 751

aggregation layer, 751

Core layer, 751

design goals, 750

availability, 751

flexibility, 751

performance, 750-751

scalability, 750

existing data center environment,
749-750

goal of, 749

on-going maintenance windows, 788

maintenance window #1, 754

port channel builds, 754-757

*port channel configuration,
757-758*

*port channel verification,
758-759*

summary, 760

maintenance window #2, 760-761

bandwidth configuration, 779

ECMP verification, 776-777

EIGRP, 776

enabling port channels, 780

FEX configuration, 786-787

HSRP routers, 777-779

*port channel configuration,
761-763*

port channel shut downs, 780

*port channel verification,
765-766*

Rapid-PVST+766

route verification, 780

routed uplink shut downs, 782

*routing table verification,
775-776*

*STP configuration, 763-764,
766*

*STP verification, 766-768,
780-781*

SVI address configuration, 783

SVI configuration, 769-774

SVI shut downs, 781-782

SVI verification, 775

*VLAN configuration, 763-764,
766*

*vPC configuration, 764-765,
768, 784-786*

migration plan, 752

premigration steps, 752-754

summary of, 788

topology of, 760, 768, 782-784

modularity (processes), 366-368

**MPLS (Multiprotocol Label
Switching), 723**

6PE, 725

6VPE, 725

CE, 719

control plane, 718

deploying, 719

device relationships, 719

enabling, 721

Forwarding table, 718

- high availability, 725-726
- IPv6, 725
- labels, 718
- Layer 3 VPN, 718, 720-723
- LDP
 - defining*, 719
 - enabling*, 721-722
 - operational commands*, 722
 - OSPF configuration*, 721-722
- LSP, 719
- LSR, 719
- managing, 725
- Nexus hardware requirements, 726-727
- NX-OS licensing, 726-727
- overview of, 717-718
- P, 719
- PE, 719
- Quality of Service, 718, 723
- RD, 719
- RSVP, 719
- RT, 719
- Traffic Engineering, 718, 723-725
 - enabling*, 724
 - IS-IS configuration*, 724
 - OSPF configuration*, 725
- troubleshooting, 725
- VRF-based CE-PE BGP configuration, 723
- MQC (Modular Quality of Service CLI)**, 644-645
- MR (Map-Resolvers)**, 730
- MS (Map-Servers)**, 730
- MSDP (Multicast Source Discovery Protocol)**, 226
 - Nexus 7000 configuration, 248-250
 - topology of, 248
- MST (Multiple Spanning Tree)**, 81, 87-91
- MTS (Message and Transaction Services)**, high availability, 366
- multicast**, 241-242
 - CoPP, 253
 - IGMP, 241-242
 - enabling*, 243
 - Nexus 1000V configuration*, 246-248
 - Nexus 1000V snooping*, 246
 - Nexus 5000 configuration*, 245
 - Nexus 7000 configuration*, 242-245
 - Nexus 7000/Nexus 5000 snooping*, 244-245
 - vEthernet1 snooping*, 246
 - multicast distribution trees, 222-223
 - multicast storms, 324
 - overview of, 221-222
- PIM, 225-226
 - ASM, 225
 - BiDir, 226, 234-235, 237
 - enabling*, 228
 - join policy configurations*, 252-253
 - MSDP, 226, 248-250
 - Nexus 7000/Nexus 5000 configuration*, 227-241
 - prune policy configurations*, 252-253
 - RP administrative scoping*, 250-251
 - SSM, 225, 239-241
 - static RPF modes*, 226, 239-241
 - topology of*, 227, 237

reverse path forwarding, 225

RP, 226-227

Multihop FCoE (FibreChannel over Ethernet), 462-463

N

NAM (Network Analysis Module)

deploying as VSB on Nexus 1010, 641-642

installing, 630-640

overview of, 629

topology of, 629

NDAC (Network Device Admission Control) and TrustSec, 282-285

NetFlow, 437-438

Nexus 1000V configuration, 442-444

Nexus 7000 configuration, 438-442

network advertisements

BGP, 194-196

OSPF configuration

Alphanumeric Process ID, 157-158

Numeric Process ID, 157

network ports (Spanning Tree Protocol), 99

network-qos policies, 646

Nexus 5x00, 664-667

Nexus 7000, 648-650

Nexus 5x00

checkpoints, 434-435

Atomic mode, 434

Best-Effort mode, 434

creating, 435-437

rollbacks, 437

Stop-at-First-Failure mode, 434

ERSPAN, 412-414

QoS, 663, 674

FEX, 668-669

forwarding architecture, 663-664

network-qos policies, 664-667

queuing policies, 667-668

SPAN, 392-393

session limits, 393-397

United Fabric configuration, 467

enabling FCoE, 468

enabling NPV mode, 468-469

FCoE NPV configuration, 469-473

Single-Hop FCoE

configuration, 469-473

verifying FCoE licenses, 467-468

Nexus 1000V

architecture of, 490-491

benefits of, 493

Control VLAN, 492-493

ERSPAN, 400-406

hypervisor, 489-490

IGMP configuration, 246-248

Installation Management Center, 510-519-525

loop prevention, 495

NAM

deploying as VSB on Nexus 1010, 641-642

installing, 630-640

overview of, 629

topology of, 629

NetFlow configuration, 442-444

overview of, 490-493

Packet VLAN, 493

- port profiles, 542-543
 - characteristics of*, 543-545
 - creating on VSM*, 545-546
 - mobility of*, 549-550
 - QoS policies*, 550-552
 - states of*, 543
 - verifying on VSM*, 546-548
- QoS, 670, 674
 - classification in*, 670-673
 - forwarding architecture*, 670
- scalability numbers, 494
- SPAN, 397-398
 - configuring*, 398-400
 - session limits*, 398
- topology of, 492
- VEM, 490, 492-494
 - functions of*, 490-491
 - installation via Installation Management Center*, 519-525
 - physical port classifications*, 495
 - port profiles*, 542-552
 - supported ports*, 494
 - VEM additions to ESX vSphere hosts*, 528-535
 - virtual port classifications*, 494-495
- VM, 493
- VMware and VMNIC, 495
- VNMC, 552-553
 - installing*, 553-562
 - PA configuration*, 570-571
 - VM-Manager*, 564-569
- VSG, 571-573
- VSM, 490, 492-493
 - changing hostnames*, 536
 - installing*, 496-541
 - Layer 3 control*, 536-542
 - PA configuration via VNMC*, 570-571
 - port profile creation*, 545-546
 - port profile verification*, 546-548
 - Telnet servers*, 536
 - uplink profile creation*, 526-528
- vSphere, 489-490
- VXLAN
 - deploying*, 604-629
 - overview of*, 602-604
 - topology of*, 603
- Nexus 1000v**
 - NX-OS licensing, 9
 - NX-OS support, 6
- Nexus 1010**
 - Nexus 1000V NAM deployments as VSB, 641-642
 - VNMC, VSG tenant creation, 579-602
 - VSG installation, 574-577
 - tenant creation in VNMC*, 579-602
 - verifying VNMC registration*, 578
 - VNMC PA configuration on VSM*, 577-578
- VSM
 - PA configuration via VNMC*, 577-578
 - verifying VNMC registration*, 578
- Nexus 2000**
 - Fabric Extenders, 60-61
 - FEX
 - Nexus 5x00 QoS*, 668-669
 - NX-OS support*, 6

- Layer 2
 - Etherchannel*, 61, 66-69
 - port-channel forwarding*, 61, 66-69
 - static pinning*, 61-66
 - Layer 2 switching, 61
 - NX-OS licensing, 9
 - Nexus 2148, NX-OS support, 6
 - Nexus 2224, NX-OS support, 6
 - Nexus 2232TM, NX-OS support, 6
 - Nexus 2232TM-E, NX-OS support, 6
 - Nexus 2232TP, NX-OS support, 6
 - Nexus 2248, NX-OS support, 6
 - Nexus 2248TP-E, NX-OS support, 6
 - Nexus 3000, NX-OS licensing, 8
 - Nexus 3016, NX-OS support, 5
 - Nexus 3048, NX-OS support, 5
 - Nexus 3064, NX-OS support, 5
 - Nexus 3548, NX-OS support, 6
 - Nexus 4000, NX-OS support, 7
 - Nexus 5000
 - IGMP configuration, 245
 - NX-OS support, 4-5
 - PIM configuration, 227-241
 - Nexus 5010
 - NX-OS support, 4
 - unsupported NX-OS features, 5
 - Nexus 5020
 - NX-OS support, 4
 - unsupported NX-OS features, 5
 - Nexus 5500, NX-OS licensing, 8
 - Nexus 5548P, NX-OS support, 4
 - Nexus 5548UP, NX-OS support, 4-5
 - Nexus 5596UP, NX-OS support, 5
- Nexus 7000
 - checkpoints, 431-432
 - Atomic mode*, 431
 - Best-Effort mode*, 431
 - creating*, 432-433
 - rollbacks*, 434
 - Stop-at-First-Failure mode*, 431
 - ERSPAN, 406-412
 - FabricPath, enabling on Nexus 7000, 122
 - IGMP configuration, 242-245
 - NetFlow configuration, 438-442
 - NX-OS licensing, 7-8
 - NX-OS support, 3-4
 - PIM configuration, 227-241
 - port channels
 - Layer 2 port-channel forwarding*, 69-71
 - load balancing options*, 103-104
 - QoS, 646, 674
 - FEX*, 661-662
 - forwarding architecture*, 646-648
 - network-qos policies*, 648-650
 - queuing policies*, 650-661
 - VDC*, 663
 - VQI*, 647-648
 - SPAN, 386-387
 - configuring*, 387-392
 - session limits*, 387
 - TrustSec configuration, 285-288
 - United Fabric configuration, 477, 484-485, 487
 - FCoE installation*, 478-479

- MDS FCoE configuration, 482-484
- Nexus 7000 shared interface allocation, 485-486
- Nexus 7000/MDS interconnection, 481-482
- QoS policies, 477-478
- Storage VDC, 478-479
- VDC configuration, 479-481
- VFC interface creation, 486-487
- VDC
 - communicating between, 42-43
 - components shared between, 38
 - configuring, 43-46
 - default VDC, 38-39
 - HA policies, 54
 - interface allocation, 46-53
 - logical segmentation on, 38-39
 - module types, 39-40
 - monitoring resources, 37-38
 - non-default VDC, 39
 - overview of, -37
 - shared interfaces, 41-42
 - storage VDC, 41
 - system resource allocation, 53-54
- Nexus 7004, NX-OS support, 3-4
- Nexus 7009, NX-OS support, 3
- Nexus 7010, NX-OS support, 3
- Nexus 7018, NX-OS support, 3
- Nexus migration case study
 - design goals, 750
 - availability, 751
 - flexibility, 751
 - performance, 750-751
 - scalability, 750
 - existing data center environment, 749-750
 - goal of, 749
 - on-going maintenance windows, 788
 - maintenance window #1, 754
 - port channel builds, 754-757
 - port channel configuration, 757-758
 - port channel verification, 758-759
 - summary, 760
 - maintenance window #2, 760-761
 - bandwidth configuration, 779
 - ECMP verification, 776-777
 - EIGRP, 776
 - enabling port channels, 780
 - FEX configuration, 786-787
 - HSRP routers, 777-779
 - port channel configuration, 761-763
 - port channel shut downs, 781
 - port channel verification, 765-766
 - Rapid-PVST+766
 - route verification, 780
 - routed uplink shut downs, 782
 - routing table verification, 775-776
 - STP configuration, 763-764, 766
 - STP verification, 766-768, 780-781

- SVI address configuration*, 783
 - SVI configuration*, 769-774
 - SVI shut downs*, 781-782
 - SVI verification*, 775
 - VLAN configuration*, 763-764, 766
 - vPC configuration*, 764-765, 768, 784-786
- migration plan, 752
- premigration steps, 752-754
- summary of, 788
- topology of, 760, 768, 782-784
- nonstop forwarding**, 370
- normal ports (Spanning Tree Protocol)**, 99
- NPIV (N-Port Identification Virtualization)**, 466
- NPV (N-Port Virtualization)**, 465-466
 - Nexus 5x00 United Fabric configuration, 468-469
 - NPIV, 466, 473-477
- NTP (Network Time Protocol)**, 444-445
- Numeric Process ID, OSPF**
 - configuring, 155-156
 - network advertisements, 157
 - passive interfaces, 159
 - stub routing, 163, 165
 - summarization, 160-161
- numeric tags**
 - EIGRP
 - instance creation*, 137-138
 - passive interfaces*, 141-142
 - stub routing*, 146
 - summarization*, 143
 - SVI, adding to EIGRP, 139-140
- Nvram files, syntax of**, 29
- NX-OS**
 - advantages of, 1-2
 - B22HP HP-FEX support, 7
 - CMP management interface, 15-18
 - continuous system operations, 2
 - CP/Supervisor management interface, 15-16
 - DCNM, 26-28
 - EXEC Command Mode, 12-13
 - FabricPath, 2
 - further reading, 57
 - Global Configuration Command Mode, 13
 - high availability, 366
 - Interface Configuration Command Modd, 13-14
 - IOS comparison, 10-12
 - licensing, 7-10
 - MPLS*, 726-727
 - VRF*, 726-727
 - management interfaces, 14-28
 - MDS 9000 support, 7
 - MGMT0 management interface, 15
 - Nexus 1000v support, 6
 - Nexus 2000 FEX support, 6
 - Nexus 3000 support, 5-6
 - Nexus 4000 support, 7
 - Nexus 5000 support, 4-5
 - Nexus 7000 support, 3-4
 - OSPF, 158
 - OTV, 2
 - PSS, 2
 - security, 2
 - SNMP management interface, 15, 23-26

SSH management interface, 15, 19-21
 supported platforms, 3-7
 Telnet management interface, 15, 18-19
 UCS support, 7
 user modes, 12-14
 VDC, 2
 vPC, 2
 XML management interface, 15, 21-22

O

on-demand diagnostics, 365
 operating system files, 35-36
 OSPF (Open Shortest Path First), 172, 725
 authentication, 167-169
 configuring, 154-155
 alphanumeric strings, 156
 network advertisements, 157-158
 Numeric Process ID, 155-156
 passive interfaces, 159-160
 Router ID, 156-157
 enabling, 155
 IOS versus NX-OS, 158
 LDP configuration as IGP in MPLS, 721-722
 MPLS TE configuration, 725
 OSPFv3 configuration, 177-178
 overview of, 154
 redistribution, 169-170
 default routes, 171-172
 Fixed Limit option, 174
 prefix list definitions, 170, 172
 redistributed route limits, 174-175
 route map definitions, 170, 172
 Warning option, 174
 Withdraw option, 174
 security, 167-169
 stub routing, 163-167
 summarization, 160-163
 topology of, 155, 170
 VRF configuration, 714

OTV (Overlay Transport Virtualization), 2, 707
 adjacency server, 692-695
 AED, 680-681
 ARP, 700-702
 control plane, 682-687
 data plane, 695-697
 encapsulation, 696
 multicast traffic, 697-698
 unicast traffic, 698
 dual adjacency, 680
 failure isolation, 698
 broadcast traffic handling, 699
 STP isolation, 698-699
 unknown unicast handling, 699
 FHRP localization, 702-704
 GLBP localization, 702-704
 HSRP localization, 702-704
 inbound path optimization, 705-707
 internal interfaces, 677-678
 join interface, 678
 LISP, 707
 multicast-enabled transport, 687-691
 multihoming, 700
 OTV Edge Device, 677, 689
 OTV site-id, 679-680

- OTV site-VLAN, 679
- overlay adjacency, 680
- overlay interface, 679
- overview of, 675-676
- QoS, 698
- site adjacency, 680
- topology of, 676
- unicast-enabled transport, 691-695
- VRRP localization, 702-704

P

P (Provider routers) and MPLS (Multiprotocol Label Switching), 719

PA (Policy-Agent)

- Nexus 1000V VSM configuration, 570-571
- Nexus 1010 VSM configuration, 577-578

packet (data) rate limits, 335-340

Packet VLAN, Nexus 1000V and, 493

PE (Provider Edge) routers and MPLS (Multiprotocol Label Switching), 719

peer templates (BGP), 192-194

PETR (Proxy ETR), 731

physical redundancy, 349-350

- cooling systems, 352-355
- EOBC, 357
- fabric modules, 357-358
- power supplies, 350-352
- supervisors, 355-357, 369-370

physical topology, 56

PIM (Protocol Independent Multicast), 225-226

- ASM, 225

- authentication, 229-230

- BiDir, 226, 234-235, 237

- enabling, 228

- join policy configurations, 252-253

- MSDP, 226

- enabling, 248-249*

- Nexus 7000 configuration, 248-250*

- Nexus 7000/Nexus 5000 configuration, 226-230

- Anycast-RP configuration, 237*

- Auto-RP configuration, 235-237*

- BiDir configuration, 234-235*

- BSR configuration, 230-232, 237*

- Hello authentication, 229-230*

- static RP configuration, 230-232*

- prune policy configurations, 252-253

- RP administrative scoping, 250-251

- security, 229-230

- SSM, 225, 239-241

- static RPF modes, 226, 239-241

- topology of, 227

PITR (Proxy ITR), 730

plain text notation, BGP configuration, 185

POAP (Power On Auto-Provisioning), 448

port channels

- ARP synchronization, 117

- flow control, 107-108

- forwarding, Layer 2 configurations, 61, 66-69, 71

- load balancing, 103-104

- load distribution, verifying, 108-109

- overview of, 103
 - physical port assignments
 - LACP*, 104-107
 - static configurations*, 104-105
 - vPC, 109-116
 - vPC-Peer Gateways*, 116
 - vPC-Peer Switches*, 116-117
 - port profiles**
 - Layer 2 configurations, 102-103
 - Nexus 1000V profiles, 542-543
 - characteristics of*, 543-545
 - creating on VSM*, 545-546
 - mobility of*, 549-550
 - QoS policies*, 550-552
 - states of*, 543
 - verifying on VSM*, 546-548
 - ports**
 - security
 - configuring*, 308-311
 - violations*, 311-313
 - STP
 - cost configuration*, 91-94
 - defining default port type*, 99-100
 - edge ports*, 99
 - network ports*, 99
 - normal ports*, 99
 - priority*, 94
 - VLAN, adding to, 73
 - power supplies, redundancy**, 350-352
 - processes**
 - modularity, 366-368
 - restarts, 368-369
 - protect action (port security violations)**, 311
 - prune policies, PIM configurations**, 252-253
 - PSS (Persistent Storage Service)**, 2, 366
 - PTP (Precision Time Protocol)**, 445-447
 - PVLAN (Private Virtual Local Area Networks)**. *See also* VLAN (Virtual Local Area Networks)
 - access switches, configuring, 78-79
 - broadcast suppression, 76
 - configuring, 77
 - defining, 78
 - enabling, 77
 - IP addresses, 76
 - overview of, 76-77
 - security, 76
 - SVI, creating, 78
 - uses for, 76
 - verifying configurations, 80
 - VLAN in
 - community VLAN*, 76-77
 - isolated VLAN*, 76
 - primary VLAN*, 76
 - secondary VLAN*, 76-77
 - Python**, 449-454
- ## Q
-
- QoS (Quality of Service)**, 674
 - ACL, 645
 - class maps, 645
 - classification matches, 645-646
 - MPLS, 718, 723
 - MQC, 644-645
 - network-qos policies, 646
 - Nexus 5x00*, 664-667
 - Nexus 7000*, 648-650

Nexus 5x00, 663, 674
FEX, 668-669
forwarding architecture,
 663-664
network-qos policies, 664-667
queuing policies, 667-668

Nexus 1000V, 670, 674
classification in, 670-673
forwarding architecture, 670

Nexus 7000, 646, 674
FEX, 661-662
forwarding architecture,
 646-648
network-qos policies, 648-650
queuing policies, 650-661
United Fabric configuration,
 477-478
VDC, 663
VQI, 647-648

OTV, 698

overview of, 643

port profiles in Nexus 1000V,
 550-552

queuing policies

Nexus 5x00 QoS, 667-668

Nexus 7000 QoS, 650-661

R

RADIUS (Remote Access Dial-In User Service)

configuration distribution, 259-264

configuring, 256-259

TrustSec and RADIUS server host
 configuration, 281-282

Rapid-PVST (Per-VLAN Spanning Tree), 80-87

rate limits (data packets), 335-340

RD (Route Distinguishers) and MPLS (Multiprotocol Label Switching), 719

redistribution

EIGRP, 149-154

OSPF, 169-177

redundancy (physical), 349-350

cooling systems, 352-355

EOBC, 357

fabric modules, 357-358

power supplies, 350-352

supervisors, 355-357, 369-370

restarts (processes), 368-369

restrict action (port security violations), 311

RLOC (Route Locator) addresses, 730, 736, 738

rollbacks (configuration)

Atomic configuration rollbacks, 33

Best-effort configuration rollbacks,
 33

checkpoints and

Nexus 5x00, 437

Nexus 7000, 434

creating, 34-35

limitations of, 34

Stop-at-first-failure configuration
 rollbacks, 33

Verbose mode configuration
 rollbacks, 33

RootGuard, 96-97

Router ID

BGP configuration, 186

OSPF configuration, 156-157

routers and MPLS (Multiprotocol Label Switching)

CE, 719

LSR, 719

P, 719
 PE, 719
 routing table verification
 (FabricPath), 126-127
 RP (Rendezvous Points)
 configuring
 Anycast-RP, 227, 237
 Auto-RP, 227, 235-237
 BSR, 226-227, 232-235
 static RP, 226, 230-232
 IP multicasts and, 226-227
 RPF (Reverse Path Forwarding), IP
 multicasting, 225-226, 239-241
 RSVP (Resource Reservation
 Protocol) and MPLS
 (Multiprotocol Label Switching),
 719
 RT (Route Targets) and MPLS
 (Multiprotocol Label Switching),
 719
 runtime diagnostics, 360-364

S

security

ACL, 302-305
 MAC ACL, 305-307
 VLAN ACL, 307-308
 BGP, 190-192
 CoPP, 327-335
 DHCP snooping, 313-316
 Dynamic ARP
 configuring, 316-317
 Inspection Trust state, 317-320
 EIGRP, 147-149
 GLBP, 215-217
 HSRP, 174-203
 IP Source Guard, 321-322

keychain management, 322-323
 LinkSec and TrustSec configuration,
 294-302
 NX-OS, 2
 OSPF, 167-169
 PIM, 229-230
 ports
 configuring, 308-311
 violations, 311-313
 PVLAN, 76
 RADIUS
 configuration distribution,
 259-264
 configuring, 256-259
 rate limits (data packets), 335-340
 security network topologies,
 255-257
 SNMPv3, 340-347
 SSH, 275-278
 TACACS+
 configuration distribution, 267
 configuring, 266-275
 traffic storm control, 323-325
 TrustSec
 AAA *configuration*, 281-294
 authentication, 279
 authorization, 279
 configuring, 278-
 LinkSec configuration, 294-302
 NDAC *definition*, 282-285
 Nexus 7000 *configuration*,
 285-288
 overview of, 278-279
 RADIUS *server host*
 configuration, 281-282
 SGACL *creation*, 292-294
 SGT *assignments*, 288-290

- SGT mapping, 290-292*
 - topology of, 279*
- unicast RPF, 325-327
- VRRP, 208-209
- VSG, 571-602
- serviceability (embedded)
 - checkpoints
 - Nexus 5x00, 431-434*
 - Nexus 7000, 431-434*
 - EEE, 447-448
 - embedded analyzer, 412-414
 - ERSPAN
 - Nexus 5x00, 412-414*
 - Nexus 1000V, 400-406*
 - Nexus 7000, 406-412*
 - NetFlow, 437-438
 - Nexus 1000V configuration, 442-444*
 - Nexus 7000 configuration, 438-442*
 - NTP, 444-445
 - POAP, 448
 - PTP, 444-445
 - Python, 449-454
 - Smart Call Home, 424-428
 - alert groups, 424*
 - configuring, 424-431*
 - executed commands, 424*
 - SPAN
 - Nexus 5x00, 392-393*
 - Nexus 1000V, 397-398*
 - Nexus 7000, 386-392*
 - overview of, 386*
- SGACL (Security Group Access Control Lists), TrustSec configuration, 292-294
- SGT (Security Group Tags), TrustSec configurations
 - SGT assignments, 288-290
 - SGT mapping, 290-292
- show commands, troubleshooting via, 54-55
- shutdown action (port security violations), 311
- Single-Hop FCoE (FiberChannel over Ethernet), 461-462, 469-473
- slot0 files, syntax of, 30
- Smart Call Home, 424-428
 - alert groups, 424
 - configuring, 424-431
 - executed commands, 424
- SNMP (Simple Network Management Protocol) management interface, 15
- SNMP management interface, 23-26
- SNMPv3 (Simple Network Management Protocol version 3), 340-347
- snooping
 - DHCP snooping, 313-316
 - IGMP
 - Nexus 1000V snooping, 246*
 - Nexus 7000/Nexus 5000 snooping, 244-245*
 - vEthernet1 snooping, 246*
- software, upgrading via ISSU (In-Service Software Upgrades), 370-383
- SPAN (Switch Port Analyzer)
 - Nexus 5x00, 392-393
 - configuring, 393-397*
 - session limits, 393*
 - Nexus 1000V, 397-398
 - configuring, 398-400*
 - session limits, 398*

- Nexus 7000, 386-387
 - configuring*, 387-392
 - session limits*, 387
- overview of, 386
- SSH (Secure Shell)**
 - configuring, 275-278
 - management interface, 15, 19-21
- SSM (Source Specific Multicast)**, 225, 239-241
- standard hosts, Layer 2
 - configurations, 101
- stateful switchovers, 369-370
- static pinning, 61-66
- static RP and RP configurations, 226, 230-232
- static RPF modes, IP multicasting and, 239-241
- static RPF (Reverse Path Forwarding), IP multicasting and, 226
- Stop-at-first-failure configuration**
 - rollbacks, 33
- Stop-at-First-Failure mode (checkpoints)**
 - Nexus 5x00, 434
 - Nexus 7000, 431
- storage
 - PSS (Persistent Storage Service), 2
 - storage VDC, 41
- Storage VDC (Virtual Device Context)**, 463-464, 478-479
- store-and-forward switching, 60
- STP (Spanning Tree Protocol)**
 - BPDU and OTV failure isolation, 698-699
 - BPDUFilter, 95-96
 - BPDUGuard, 94-95
 - Bridge Assurance, 98-99
 - dispute mechanism, 98
 - errdisable recovery, 95
 - LoopGuard, 97-98
 - MST, 81, 87-91
 - overview of, 80
 - port costs, configuring, 91-94
 - port priority, configuring, 94
 - ports
 - defining default port type*, 99-100
 - edge ports*, 99
 - network ports*, 99
 - normal ports*, 99
 - Rapid-PVST, 80-87
 - RootGuard, 96-97
 - timers, 84-87
 - TrunkFast, 100
 - virtualization hosts, 100-103
 - VLAN, 81
- Strict Unicast RPF mode (unicast RPF)**, 325
- stub routing**
 - EIGRP, 145-147
 - OSPF, 163-167
- summarization and EIGRP (Enhanced Interior Gateway Routing Protocol)**, 142-145
- supervisors, redundancy**, 355-357, 369-370
- SVI (Switched Virtual Interfaces)**
 - EIGRP, adding to, 139-141
 - GLBP configuration, 213
 - HSRP configuration, 200
 - PVLAN, creating for, 78
 - VRRP configuration, 206-207
- switch IDs, configuring via FabricPath**, 122

switching

- cut-through switching, 60

- Layer 2 switching, Nexus 2000 and, 61

- store-and-forward switching, 60

switchovers (stateful), 369-370**system files**

- management, 28

- configuration files, 33-35*

- file systems, 28-33*

- operating system files, 35-36*

- syntax of, 30

System Manager (NX-OS), 368-369

- high availability, 366

- process restarts, 368-369

- stateful switchovers, 368-369

T**TACACS+ (Terminal Access Controller Access-Control System+)**

- configuration distribution, 267

- configuring, 266-267

- global TACACS+ keys, 268*

- TACACS+ server groups, 269-270*

- TACACS+ server hosts, 268-269*

- TACACS+ server keys, 268*

- TACACS+ source interface, 270-275*

- enabling, 266-267

Telnet management interface, 15, 18-19**Telnet servers, VSM and, 536****templates (peer) and BGP (Border Gateway Protocol), 192-194****timers (Spanning Tree Protocol), 84-87****topologies, 56, 237**

- BGP, 184

- FabricPath, 122-123

- FCoE, Nexus 7000/MDS topology, 481

- GLBP, 212

- IGMP

- Nexus 1000V configuration, 247*

- Nexus 7000 configuration, 242*

- MSDP, 248

- OSPF, 155, 170

- OTV, 676

- OTV adjacency server, 692-693

- physical topology, 56

- PIM, 227, 237

- security network topologies, 255-257

- TrustSec, 279

- VRRP, 205

Traffic Engineering and MPLS (Multiprotocol Label Switching), 718, 723-725

- enabling in, 724

- IS-IS configuration, 724

- OSPF configuration, 725

traffic storm control, 323-325**troubleshooting, 54**

- debug commands, 54-56

- MPLS, 725

- show commands, 54-55

trunk ports, Layer 2 configurations, 100-101**TrustSec**

- authentication, 279

- authorization, 279

configuring, 278-302
 AAA configuration, 281-294
 NDAC definition, 282-285
 Nexus 7000 configuration, 285-288
 RADIUS server host configuration, 281-282
 SGACL creation, 292-294
 SGT assignments, 288-290
 SGT mapping, 290-292
 LinkSec configuration, 294-302
 overview of, 278-279
 topology of, 279

U

UCS (Unified Computing Systems), NX-OS support, 7

UDLD (Unilateral Link Detection), 118-119

unicast

unicast RPF, 325-327
 unicast storms and traffic storm control, 324-325

Unified Fabric

benefits of, 456, 488
 Ethernet, 456
 10-Gigabit Ethernet, 456-458
 FCoE, 458-460, 467-468
 FCoE NPV mode, 466, 473-477
 Multihop FCoE, 462-463
 Single-Hop FCoE, 461-462, 469-473
 Nexus 5x00 configuration, 467
 enabling FCoE, 468
 enabling NPV mode, 468-469
 FCoE NPV configuration, 469-473

Single-Hop FCoE configuration, 469-473
verifying FCoE licenses, 467-468

Nexus 7000 configuration, 477, 484-485, 487
 FCoE installation, 478-479
 MDS FCoE configuration, 482-484
 Nexus 7000 shared interface allocation, 485-486
 Nexus 7000/MDS interconnection, 481-482
 QoS policies, 477-478
 Storage VDC, 478-479
 VDC configuration, 479-481
 VFC interface creation, 486-487

NPV, 465-466
 FCoE NPV mode, 466, 473-477
 NPIV, 466

overview of, 455-456
 storage VDC, 463-464

upgrades, ISSU (In-Service Software Upgrades), 370-383

uplink profiles, creating for Nexus 1000V VSM, 526-528

usb1 files, syntax of, 29

usb2 files, syntax of, 30

user modes, 12-14

V

VDC (Virtual Device Contexts), 2, 663

communicating between, 42-43
 components shared between, 38
 configuring, 43-46

- default VDC, 38-39
- HA policies, 54
- interface allocation, 46-53
- module types, 39-40
- Nexus 7000
 - logical segmentation on*, 38-39
 - QoS, 663
 - Storage VDC*, 478-479
 - VDC configuration*, 479-481
- non-default VDC, 39
- overview of, 37
- resources, monitoring, 37-38
- shared interfaces, 41-42
- storage VDC, 41, 463-464
- system resource allocation, 53-54
- VEM (Virtual Ethernet Module)**
 - Nexus 1000V, 490, 492-494
 - port profiles*, 542-552
 - VEM additions to ESX vSphere hosts*, 528-535
 - VEM functions in*, 490-491
 - VEM installation via Installation Management Center*, 519-525
 - VEM physical port classifications*, 495
 - VEM supported ports*, 494
 - VEM virtual port classifications*, 494-495
- Verbose mode configuration**
 - rollbacks, 33
- vEth (virtual Ethernet) ports, VEM in Nexus 1000V**, 494
- vEthernet1, IGMP configuration**, 246-248
- VFC interface**
 - Nexus 5x00 United Fabric configuration
 - FCoE NPV configuration*, 476-477
 - Single-Hop FCoE configuration*, 472-473
 - Nexus 7000 United Fabric configuration, 486-487
- virtual machine managers. See hypervisor**
- Virtual NIC, VEM in Nexus 1000V**, 494
- virtualization hosts**
 - linking to, 101
 - STP, 100-103
- VLAN (Virtual Local Area Networks). See also PVLAN (Private Virtual Local Area Networks)**
 - community VLAN, 76-77
 - configurations, verifying, 74-76
 - configuring, 72
 - creating, 72
 - FabricPath, configuring VLAN via, 123-124
 - internal VLAN, displaying, 71-72
 - isolated VLAN, 76
 - memberships, assigning, 73-74
 - multiple VLAN, creating, 72
 - Nexus 5x00 United Fabric configuration
 - FCoE NPV configuration*, 474-475
 - Single-Hop FCoE configuration*, 471-472
 - OTV site-VLAN, 679
 - overview of, 71
 - ports, adding, 73
 - PVLAN
 - community VLAN*, 76-77
 - isolated VLAN*, 76

- primary VLAN*, 76
- secondary VLAN*, 76-77
- STP, 81
- trunk interfaces, configuring, 74
- trunks, adding/removing from, 74
- VTP, 72-73
- VLAN ACL (Access Control Lists)**, 307-308
- VM (Virtual Machines) and Nexus 1000V**, 493
- vmknic (virtual kernel NIC), VEM in Nexus 1000V, 494
- VM-Manager and Nexus 1000V VNCM**, 564-569
- VMNIC (Virtual Machine NIC) and VMware**, 495
- VMware**, 490, 491
 - VMNIC, 495
 - VSM uplink profile creation, 526-528
- vNIC (virtual NIC), VEM in Nexus 1000V**, 494
- VNMC (Virtual Network Management Center), Nexus 1000V and**, 552-553
 - installing, 553-562
 - PA configuration, 570-571
 - VM-Manager, 564-569
 - VSG tenant creation, 579-602
- VNP ports, Nexus 5x00 United Fabric configuration**, 475-476
- volatile files, syntax of, 30
- vPC (Virtual Port Channels)**, 2, 109-116
 - HSRP/VRRP interactions, 210-211
 - vPC-Peer Gateways, 116
 - vPC-Peer Switches, 116-117
- vPC+ (Virtual Port Channel+)**, 127-132
- VPN (Virtual Private Networks), Layer 3 VPN and MPLS**, 718, 720-723
- VQI (Virtual Queuing Index), Nexus 7000 QoS**, 647-648
- VRF (Virtual Routing and Forwarding)**
 - EIGRP configuration, 714-715
 - Interface and VRF configuration, 712
 - IS-IS configuration, 715
 - MPLS, VRF-based CE-PE BGP configuration, 723
 - Nexus hardware requirements, 726-727
 - NX-OS licensing, 726-727
 - operational commands, 713
 - OSPF configuration, 714
 - overview of, 709
 - predefined VRF, 710-713
 - routing table isolation, 712-713
 - VRF-capable components, 710
 - VRF-Lite, 713-717
 - route import/export functionality*, 716
 - route import/export verification*, 716-717
- VRRP (Virtual Router Redundancy Protocol)**, 198
 - authentication, 208-209
 - configuring, 205-208
 - enabling, 206
 - OTV VRRP localization, 702-704
 - overview of, 205
 - priority and preempt, 207
 - secondary IP addresses/subnets, 209
 - security, 208-209

- topology of, 205
- vPC interactions, 210-211
- VSAN (Virtual Storage Area Network) mapping, Nexus 5x00 United Fabric configuration**
 - FCoE NPV configuration, 474-475
 - Single-Hop FCoE configuration, 471-472
- VSB (Virtual Service Blades), Nexus 1000V deployments, 497-509, 641-642**
- VSG (Virtual Security Gateways)**
 - Nexus 1000V configuration, 571-573
 - Nexus 1010 installation, 574-577
 - tenant creation in VNMC, 579-602*
 - verifying VNMC registration, 578*
 - VNMC PA configuration on VSM, 577-578*
- VSM (Virtual Supervisor Mode)**
 - Nexus 1000V, 490, 492-493
 - changing hostnames, 536*
 - installing on, 496-541*
 - Layer 3 control, 536-542*
 - PA configuration via VNMC, 570-571*
 - port profile creation, 545-546*
 - port profile verification, 546-548*
 - Telnet servers, 536*
 - uplink profile creation, 526-528*
 - Nexus 1010
 - PA configuration via VNMC, 577-578*
 - verifying VNMC registration, 578*

- vSphere and Nexus 1000V, 489-490, 528-535
- vswwif, VEM in Nexus 1000V, 494
- VTP (VLAN Trunking Protocol), 72-73**
- VXLAN (Virtual Extensible LAN)**
 - deploying, 604-629
 - overview of, 602-604
 - topology of, 603

W

Warning option

- EIGRP redistribution, 152
- OSPF redistribution, 174
- web resources, NX-OS, 57
- Withdraw option, 174**
 - EIGRP redistribution, 152-153
 - OSPF redistribution, 174

X - Y - Z

- XML management interface, 15, 21-22
- xTR (ITR and ETR), 730