




SDN and NFV Simplified

A Visual Guide to Understanding Software Defined Networks and Network Function Virtualization



Jim Doherty

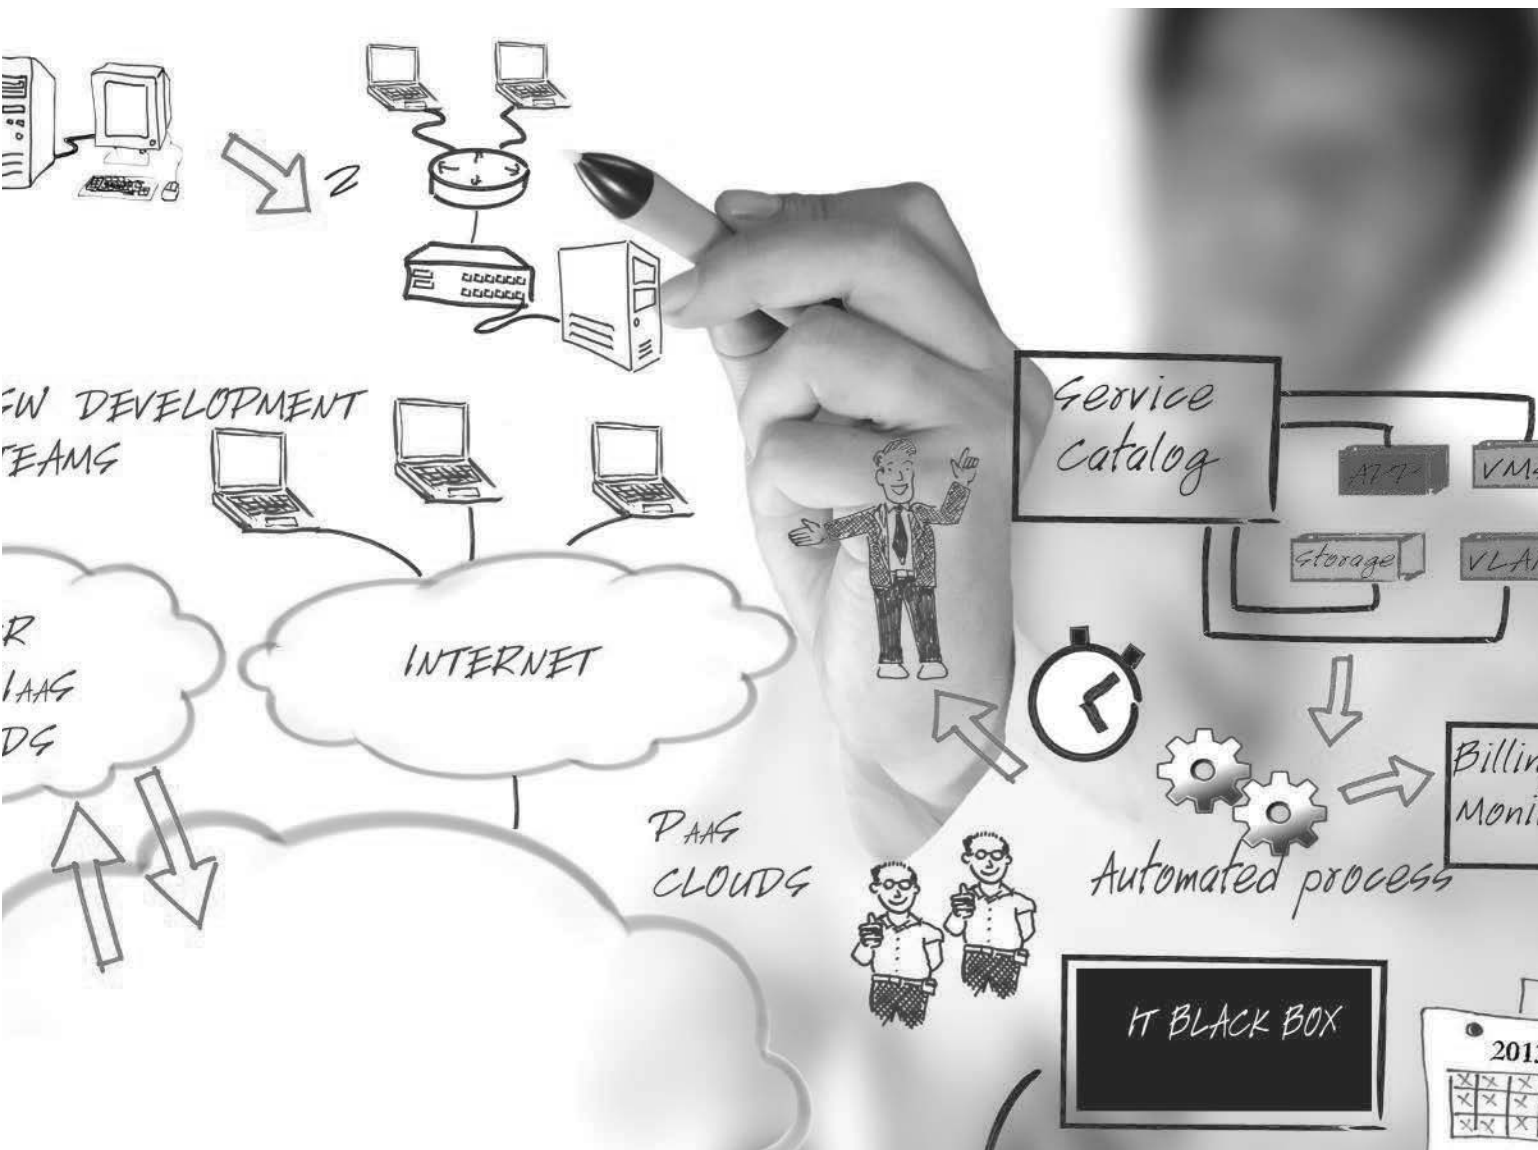
FREE SAMPLE CHAPTER



SHARE WITH OTHERS

SDN and NFV Simplified

A Visual Guide to Understanding Software Defined
Networks and Network Function Virtualization



SDN and NFV Simplified

Jim Doherty

Copyright © 2016 Pearson Education, Inc.

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

ISBN-13: 978-0-13-430640-7

ISBN-10: 0-13-430640-6

Library of Congress Control Number: 2015956324

Text printed in the United States on recycled paper at RR Donnelley in Kendallville, Indiana.

First printing: March 2016

Many of the designations used by manufacturers and sellers to distinguish their products are claimed as trademarks. Where those designations appear in this book, and the publisher was aware of a trademark claim, the designations have been printed with initial capital letters or in all capitals.

The author and publisher have taken care in the preparation of this book, but make no expressed or implied warranty of any kind and assume no responsibility for errors or omissions. No liability is assumed for incidental or consequential damages in connection with or arising out of the use of the information or programs contained herein.

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

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

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

Visit us on the Web: informit.com/aw

Editor-in-Chief

Dave Dusthimer

Executive Editor

Mary Beth Ray

Development Editor

Jeff Riley

Managing Editor

Sandra Schroeder

Project Editor

Mandie Frank

Copy Editor

Keith Cline

Indexer

Tim Wright

Proofreader

Debbie Williams

Technical Reviewer

Brian Gracely

Editorial Assistant

Vanessa Evans

Designer

Mark Shirar

Composer

Studio Galou

To Katie, Samantha, and Conor

Contents

Introduction xiii

Part 1 Virtualization 101: The Basics of Virtualization 1

Chapter 1 Primer on Virtualization 3

- Server Proliferation, Massive Power Bills, and Other IT Nightmares 3
- How Servers Work 6
- How VMs Fix the Underutilized Server Problem 7
- Enter the Hypervisor 8
- Why Are Virtual Machines Such a Big Deal? 10

Chapter 2 Benefits of Virtual Machines 13

- Reduced Cost 13
- Less Space (Even More Cost Savings) 15
- Availability and Flexibility 15
- Faster Application Spin-Up and Provisioning 16
- Easier Access for Development 17
- Believe the Hype! 18

Chapter 3 Hypervisors (VMWare, KVM, and Others) 21

- An Operating System for Operating Systems 21
- A Virtual Machine Monitor 22
- Types of Hypervisors 22
- Hypervisor Vendors 24
 - KVM 24
 - Xen 25
 - VMware ESXi 26
 - Microsoft Hyper-V 26
- Choosing a Hypervisor 27
- Summary 27

Chapter 4 Managing Virtual Resources 29

- What Is a Workload? 30
- Managing Virtual Resources in the Hypervisor 31
- Virtual Resource Providers and Consumers 33
- So How Do You Manage Virtual Resources? 34

Part 2 Virtualization 201: Virtualizing the Data Center (a.k.a. Clouds) 37

Chapter 5 Virtualized Data Centers (Some Call Them Clouds) 39

- Benefits of Virtualizing the Data Center 39
 - Less Heat Buildup 39
 - Reduced Hardware Spend 40
 - Faster Deployment 40
 - Testing and Development 40

	Faster Redeploy	40
	Easier Backups	40
	Disaster Recovery	41
	Server Standardization	41
	Separation of Services	41
	Easier Migration to the Cloud	41
	Is It a Cloud Yet?	41
	The Five Cloud Attributes	42
	On-Demand Self-Service	43
	Ubiquitous Network Access	43
	Pay Per Use	43
	Rapid Elasticity	43
	Location-Independent Resource Pooling	44
	Types of Clouds	44
	Software as a Service	44
	Infrastructure as a Service	45
	Platform as a Service	47
	Cloud Deployment Models	48
	Private Clouds	48
	Shared Multitenant Clouds	48
	Public Clouds	48
	Hybrid Clouds	48
Chapter 6	Virtual Machine Connectivity	53
	Networking in Traditional Data Centers	53
	Virtualized Data Center Design	55
	Addressing with Virtual Machines	56
Chapter 7	Networking Gear in Virtualized Data Centers	61
	The Evolution of Data Center Switching	61
	Cloud and Data Center Layout and Architecture	63
	Virtualized Aware Network Switches	65
Chapter 8	VMware, vSphere, VMotion, and VXLAN	67
	VMware Product Design	67
	vSphere	68
	VMotion	69
	VXLAN	70
	VXLAN Tunnel Endpoints	71
	Summary	72
Chapter 9	Multitenancy and the Problems of Communal Living	73
	SaaS Multitenancy	73
	Pros and Cons of SaaS Multitenancy	75
	IaaS Multitenancy	76
	Pros and Cons of IaaS Multitenancy	77

Part 3 Network Functions Virtualized: Why Stop With Servers? 81

Chapter 10 How Do You Virtualize a Network? 83

- Network Virtualization 83
- How Does This Fit with NFV and SDN? 84
 - Server Virtualization 85
 - Network Virtualization 85
 - Network Functions Virtualization 85
 - Software-Defined Networking 85
- Virtualizing the Network 86

Chapter 11 Virtualizing Appliances 89

- Layer 4 Through 7 Network Services 89
 - Firewalls 90
 - VPNs 90
 - SSL Offload 90
 - Load Balancer 90
- Fighting Virtualization with Virtualization 91
- What's the "So What"? 92

Chapter 12 Virtualizing Core Networking Functions 93

- Virtualization Recap 93
- Where Core Functions Are Being Virtualized 95

Chapter 13 What About Scalability and Performance? 99

- Scalability Versus Performance 99
- Performance in Network Virtualization 100
- Scalability and Performance in Virtual Networks 100
- Scalability and Performance for Virtual Appliances 101
- Scalability and Performance of Virtualized Networks 102
- Summary 102

Part 4 Modern Networking Approaches to Virtualization 105

Chapter 14 From Consumers to Creators 107

- The Emergence of SaaS 108
- Cloud Business Consumer-Creators 109

Chapter 15 OpenFlow 113

- OpenFlow History 114
- How OpenFlow Works 115

Chapter 16 VMware Nicira 119

- VMware NSX 120
 - Network Virtualization with NSX 120
 - How VMware Leverages Nicira (NSX) 121

Chapter 17	Cisco Insieme	125
	Cisco’s Hybrid SDN Solution	125
	Cisco SDN and Insieme	126
Chapter 18	OpenStack	129
	Applications on Modern Networks	129
Part 5	Software Defined Networks	135
Chapter 19	The Evolution of the Data Center Network	137
	Networks Worked Great, Until They Didn’t	138
	Traditional Data Center Design Goals	139
	High Availability	139
	Low Latency	139
	Scalability	139
	Security	139
	The Cost Model Explodes	140
	How We Got Here	141
Chapter 20	What’s Wrong with the Network We Have?	143
	A Brief Review of Networking	144
	Control Planes and Forwarding Planes	145
	The Cost of Complexity	145
	Decoupling Networking Applications from Networking Gear	147
Chapter 21	How SDN Works	149
	Understanding SDN	149
	The Application Layer	151
	The Control Layer	152
	The Infrastructure Layer	152
	A Programmable Network	152
	So What’s the “So What?”	153
Chapter 22	The Economic Impact of SDN, NFV, and the Cloud	157
	Winners in SDN, NFV, and the Cloud	157
	How the “Little Guy” Wins	157
	How Large Enterprises Win with SDN, NFV, and the Cloud	160
	Losers in the Cloud	163
	The Economic Value of Increased Innovation	164
Part 6	SDN Controllers	167
Chapter 23	SDN Controllers	169
	Centralized Control	169
	Commercial Versus Open Source Controllers	170
	Network Virtualization	171

Chapter 24	The OpenDaylight Project	175
	How the ODL Architecture Works	177
	The ODL Controller Platform	178
Chapter 25	The Fight to Control Your Network	181
	Separation of Internal Controls	181
	You Can See It, But Who Controls It?	184
Chapter 26	What's the Business Case for SDN?	187
	SDN Use Case Examples	188
	Data Center Optimization	189
	Network Access Control	189
	Network Virtualization	189
	Virtual Customer Edge	189
	Dynamic Interconnects	189
	Virtual Core and Aggregation	189
	Summary	190
Part 7	Virtualized Networks	193
Chapter 27	Goodbye Truck Rolls	195
	Data Center Scale	195
	A New Maintenance Philosophy	198
	Summary	199
Chapter 28	What If the Shoe Doesn't Fit?	201
	Where SDN Does Not Fit	202
	When Should You Adopt SDN?	202
	Stuck in the Middle	203
Chapter 29	Service Chaining	205
	Service Chaining in SDN	206
Chapter 30	NFV: What Happens to All the Network Appliances?	209
	How Network Appliances Are Different	209
	Replacing Big Hardware Appliances with Many Small Virtual Appliances	210
	When Not to Get Rid of an Appliance	211
Part 8	Security	213
Chapter 31	Where's My Data, Exactly?	215
	Storage Virtualization	215
	Storage-Area Networks	216
	Data Location and Security	217
	So What Are the Nontechnical Issues That We Need to Address?	218
	Summary	219

Chapter 32	Preventing Data Leakage	223
	Minimizing Data Loss	224
	Data Loss Prevention	225
Chapter 33	Logging and Auditing	229
	Where Logging Matters	231
	Summary	233
Chapter 34	Encryption in Virtual Networks	235
	Data in Motion	235
	Data at Rest	236
	Key Management	238
	Best Practices	238
Chapter 35	Everything Old Is Now New Again	241
	How We Got Here	241
	The Mainframe Model	241
	The Personal Computer Model	241
	The Networked Model	242
	The Internet Model	242
	Grid Computing Model	242
	Cloud Computing Model	242
	What We Have Learned	242
	Retro Security Considerations	244
	Recycled Ideas on Mobile and Web Apps	245
Part 9	Visibility	247
Chapter 36	Overlay Networks	249
	MPLS: The Original Virtual Network	249
	Virtual Layer 2 Designs	250
	Enter SDN	252
	Common Encapsulation Techniques	252
Chapter 37	Network Management Tools	255
	What's in the Tool Bag?	256
	Tapping In	256
	Gaining Visibility	257
Chapter 38	Quality of Experience	261
	Deep Packet Inspection	263
Chapter 39	Monitoring Traffic Between Virtual Switches	265
	Getting VM Visibility	265
	Monitoring VM-to-VM traffic	266
	How VxLANs Work	267
	Creating a "Visibility Layer"	267

Part 10 The Big Picture 269

Chapter 40 Pulling It All Together 271

- Why the Network Had to Change 271
- How SDN and NFV Tie Together 273
- SDN's Downside: A Loss of Visibility 274
- SDN Orchestration 274

Chapter 41 How SDN and NFV Will Affect You 277

- Operational Domains 278
 - Mobility Virtualization 278
 - Virtual CPE and Service Chaining 278
 - NFV and Service Orchestration 278
 - WAN Optimization and Innovation 278
 - Network Optimization 278
 - Policy-Driven Application Provisioning and Delivery 279
- SDN Use Cases 279
 - Network Access Control 279
 - Network Virtualization 279
 - Data Center Optimization 279
 - Direct Inter-Connects 279
- Embracing SDN and NFV 279

Chapter 42 What's Next in Networking? 283

- Separate but Complementary 283
- Virtual Customer Premise Equipment 284
 - SDN and NFV Working Together 285
- Summary 286

Index 287

Acknowledgments

I would like to thank the following people.

First and foremost, I want to thank my researcher and writing assistant, Alasdair Gilchrist. This book simply could not have happened without his hard work and dedication.

This book has been in the works for quite a while, and over the course of a few years, it has morphed from what was originally called *Cloud Networking Simplified*. For a variety of reasons, there were stops and starts, do overs, and evolutions, and in a couple of cases, potential co-authors were brought in. Among them were Dave Asprey, who was a very knowledgeable tech guy before becoming the “Bullet Proof Coffee” guy, and Brian Gracely, who is still a very knowledgeable tech guy who ended up being the technical reviewer of this version of the book. While this book is its own thing, several of their original contributions influenced this version.

Brian Gracely also gets his own thank you for being the technical reviewer of this book. He kept me on my toes and spent a lot of time reading, reviewing, and verifying my work. This book is much better than it would have been without his efforts.

Finally, I want to thank the great team at Pearson. Always the consummate professionals, they make my job much easier, and they are all a pleasure to work with. Specifically, I’d like to thank Executive Editor Mary Beth Ray, who had the patience to keep this project alive through all the iterations. Development Editor Jeff Riley for being a great “man behind the curtain.” Line Artist Laura Robbins for helping with the whiteboard layouts and line art. And last but not least, the editing team, including Project Editor Mandie Frank, Development Editors Drew Cupp and Jeff Riley, and Copy Editor Keith Cline, who I’m sure I drove crazy with my many spelling errors, typing errors, run-on sentences, and use of passive voice. This is my sixth book with Pearson, and every crew I’ve worked with has been outstanding. This one is no exception. Thank you all.

Jim Doherty

About the Author

Jim Doherty has more than 17 years of engineering and marketing experience across a broad range of networking, security, and technology companies. Focusing on technology strategy, product positioning, and marketing execution, Jim has held leadership positions for Cisco Systems, Certes Networks, Ixia, and Ericsson Mobile. Currently, he is the SVP of Sales and Marketing for Percona.

Doherty is also the creator and co-author of the *Networking Simplified* series of books, which includes *Cisco Networking Simplified*, *Home Networking Simplified*, and several other titles. He has also written books on mobile security and other networking topics.

Jim is a former U.S. Marine Corps Sergeant and holds a Bachelor of Science degree in electrical engineering from North Carolina State University and an MBA from Duke University. Jim lives in Raleigh, North Carolina, with his wife and two children.

Introduction

Dear Reader,

Thank you for buying (or considering) this book. By way of background, this book is the latest in a series of *Networking Simplified* books that have covered topics such as basic networking with *Cisco Networking Simplified*, and other titles that have taken a “simplified” approach to Voice over IP (VoIP), security, and home networking.

If you are not familiar with the approach of the *Networking Simplified* books, the idea is to explain networking concepts and technologies to people who are not intending to become technical experts, but want to know what underlying networking technologies are and how they work at a high level. This includes business and marketing folks, salespeople, investors, people who work in technical companies but in nontechnical roles, and even technical people from different vocations who want a cursory introduction to networking.

In others words, I wanted to write a book that explains how all this stuff works, and why it matters, without assuming people needed to configure a router or set up a wireless network, and without assuming the readers were dummies. The aim is to not go too deep but not be too light either.

To that end, I’ve now written (with the help of many people listed in the acknowledgments) a book on the next big wave in networking: software-defined networking (SDN) and network functions virtualization (NFV). My approach to this book is similar to the previous books in that I start with some foundational topics, which in this case is virtualization. In fact, the first several sections of the book focus on this because 1) it’s really important and 2) I did not think that there was a good resource out there that took the same simplified approach. The book then moves into SDN and NFV, and through it all are parts about cloud networks, virtualized data centers, and network virtualization.

Whiteboards

Also throughout this book, typically at the end of chapters, are “whiteboard” diagrams. These hand-drawn (and then digitized) whiteboards are meant to capture key aspects of the chapter topics and present those ideas as if they were drawn on a whiteboard.

In keeping with the at-a-glance feature of the previous *Networking Simplified* books, the whiteboards are meant as enhancements to the text, and it is hoped that they will help aid in your understanding of the topics they cover.

Who Is This Book For, and What Will You Get Out of It?

This book is mostly meant for nontechnical people who want to know what SDN, NFV, virtualization, and cloud networking are about: what they are, how they all work, and why they are important. You won’t walk away from this book knowing how to configure an SDN controller, but you will know what one is and what it does. If you are a technical person with no exposure to these topics, this book can serve as a high-level introduction.

How This Book is Organized

I've broken the book into ten sections, each covering a central theme. This allows you to bounce around and pick one topic at time. If this is all brand new to you, tackling the topics in order is recommended, because they tend to build on each other.

Part 1: Virtualization 101: The Basics of Virtualization

As the name implies, Part I provides a primer on virtualization and virtual machines (VMs): what they are, how they work, and why they were adopted so quickly.

Part 2: Virtualization 201: Virtualizing the Data Center (a.k.a. Clouds)

Once you understand what VMs are, you need to know how they work together and what happens when you have an entire data center of them. This part focuses on the particulars of how clouds work and how they are managed.

Part 3: Network Functions Virtualized: Why Stop With Servers?

Servers aren't the only thing you can virtualize. In fact, many of the traditional applications such as firewalls can be virtualized, too. NFV is a key aspect of modern networking. This part explores not only the technology, but also the impact on how the networks work as a result.

Part 4: Modern Networking Approaches to Virtualization

Understanding how virtualization works is a great start, but unless you can connect to them over a network, a VM is not very useful. It turns out, however, that connecting to VMs requires some changes to networking before you can leverage their full power. This part explores some of the ways modern networks are managed.

Part 5: Software Defined Networks: Virtualizing the Network

SDN represent a pivot point in how networks are built and who manages them. The technology is analogous to the virtualization of data centers, and in fact, the shift to SDN is in large part an attempt to leverage the most value out of virtualized data centers and clouds. This part covers the basics of SDN technology and the economic impact of its adoption.

Part 6: SDN Controllers

Furthering the SDN discussion, this part goes into depth about SDN controllers. There's a lot of competition among many companies to be among the leaders in this technology, and given the economic benefits of being among the leaders, the stakes are high. In addition to looking at specific controllers, this part also looks at the economics of the technology.

Part 7: Virtualized Networks: Connecting It All Together

With an understanding of the pieces, this part puts virtualization, SDN, and NFV together. More than just how it works, the chapters here explore building and managing virtualized networks.

Part 8: Security: The Security Thing

Security is an ever-present part of networking these days, and in virtualized networks, where you don't always know where your data assets are, security is always an interesting topic. The chapters in this part deal with the unique security aspects of virtualization and virtualized networks.

Part 9: Visibility

You can only manage and control what you see, and this is no different in virtualized networks. The problem is that the very nature of virtualization creates blind spots in the network because data and information can move in a way that traditional network monitoring cannot detect. This part looks at how virtual networks are monitored.

Part 10: The Big Picture

In the last part of the book, all the topics previously mentioned come together for a view of how it all works together. Here we also look at how this technology will affect the way people work and communicate, and we take a peek into what might be coming next.

In the end, I'm hopeful that you will find this book to be useful, informative, and interesting. I sincerely appreciate having the opportunity to share this information with you.

—Jim Doherty

This page intentionally left blank

This page intentionally left blank



How Do You Virtualize a Network?

This chapter attempts to answer an important question: How do you virtualize a network? Before attempting an answer, though, it's a good time to take a step back and answer a couple of “big” questions, such as:

- What exactly is network virtualization, and how does it relate to the virtualization covered so far in this book?
- How does network virtualization fit into the grand scheme of network functions virtualization (NFV) and software-defined networking (SDN)?

Once we answer these questions, it's much easier to answer the question that the chapter title poses. More importantly, these answers provide the framework of *why* we would want to virtualize the network.

Network Virtualization

As mentioned throughout the first two sections of this book, virtualization (which typically means server virtualization when used as a standalone phrase) refers to the abstraction of the application and operating system from the hardware.

In a similar way, network virtualization is the abstraction of the network endpoints from the physical arrangement of the network. In other words, network virtualization allows you to group or arrange endpoints on a network independent from their physical location.

It's worth noting that network virtualization is nothing new. In fact, it's been around a long time. The most common forms of network virtualization are virtual LANs (VLANs), virtual private networks (VPNs), and Multiprotocol Label Switching (MPLS). All of these technologies essentially enable the administrators to group physically separate endpoints into logical groups, which makes them behave (and appear) as if they are all on the same local (physical) segment. The ability to do this allows for much greater efficiencies in traffic control, security, and management of the network.

In many cases, this type of virtualization is performed via some form of encapsulation whereby messages or traffic between endpoints in the same logical group are “packaged” into another message that is better suited for transport over a physical segment of the network. Once the message has reached the endpoint, the original message is unpacked, and the intended endpoint receives the message in the same format as it would have if the two endpoints were on the same physical segment of the network.

Figure 10-1 illustrates one way that VLANs would be used. In this case, workers in different departments work on multiple floors of a building. A single switch can service each floor of the building, such that all workers on a given floor would be part of the same network segment. VLANs allow you to logically group endpoints so that they all look as if they are on the same segment. Further, this can be done across many buildings or even across large networks where endpoints are scattered all over the globe—although care should be taken when extending VLANs over long distances because they can create fragile networks.

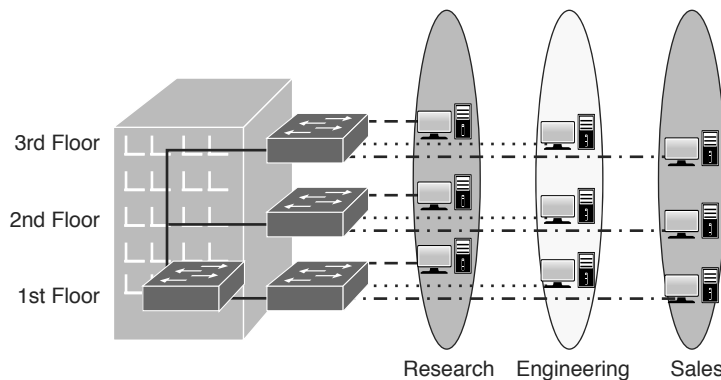


Figure 10-1 *Virtual LANs, or VLANs, were an early form of network virtualization that allowed physically separate endpoints to behave as if they were all connected to the same local switch.*

It turns out that this good old technique that has been around for many years makes server virtualization, or more accurately connecting VMs, much easier and much more efficient. It’s easy to see why when you imagine the VMs being spun up here, there, and everywhere in a virtualized data center or cloud, and then being paused, moved, started again, or even being moved while still being active.

With all that spontaneous creation and movement that is done without any regard for the specific physical location in the data center (or even with regard to a specific data center), having the ability to create and manage logical groupings becomes critical.

How Does This Fit with NFV and SDN?

With a basic grasp of what server virtualization is from previous chapters and with the newly gained understanding of network virtualization, it’s worth spending a few words on how they are related to network functions virtualization (NFV) and software-defined networking (SDN). To keep it in context, all four topics are summarized here.

Server Virtualization

Server virtualization is the abstraction of applications and operating systems from physical servers. This allows for the creation of VMs (app and OS pairs) that offer much greater usage efficiency on physical servers and afford enormous flexibility with regard to provisioning of applications.

Network Virtualization

Network Virtualization refers to the creation of logical groupings of endpoints on a network. In this case, the endpoints are abstracted from their physical locations so that VMs (and other assets) can look, behave, and be managed as if they are all on the same physical segment of the network. This is an older technology, but one that is critical in virtual environments where assets are created and moved around without much regard for the physical location. What is new here is the automation and management tools that have been purposely built for the scale and elasticity of virtualized data centers and clouds.

Network Functions Virtualization

NFV refers to the virtualization of Layer 4 through 7 services such as load balancing and firewalling. Basically, this is converting certain types of network appliances into VMs, which can then be quickly and easily deployed where they are needed. NFV came about because of the inefficiencies that were created by virtualization. This is a new concept; so far, only the benefits of virtualization have been covered, but virtualization causes a lot of problems, too. One of them was the routing of traffic to and from network appliances that typically were located at the edge of the data center network. With VMs springing up and being moved all over, the traffic flows became highly varied, which caused problems for fixed appliances that had to serve the traffic. NFV allows us to create a virtual instance of a function such as a firewall, which can be easily “spun up” and placed where it is needed, just as they would a VM. Much of this section focused on how this is done.

Software-Defined Networking

SDN refers to the ability to program the network. SDN is a newer technology, one that was born as a result of virtualization and the shift of where the “chokepoint” is in data communications. In short, the ability to set up or make changes to a network cannot keep up with the ability to provision applications with a click of a button. SDN makes the network programmable (which means network admins can quickly make adjustments to the network based on changing requirements). SDN is made possible by separating the control plane (the brains of the network) from the data plane (the muscle of the network). SDN is covered in depth in Part 5, “Software Defined Networks: Virtualizing the Network,” and Part 6, “SDN Controllers,” of this book.

All four of these technologies are designed to improve the mobility, agility, and flexibility of networks and data communication. However, virtualization, network virtualization, and network functions virtualization can all work on existing networks because they reside on servers and interact with “groomed” traffic sent to them. SDN, however, requires a new

network topology and SDN-aware devices where the data and control planes are separate and programmable.

Virtualizing the Network

One of the reasons it's a good idea to make the change to network virtualization is that it allows network admins and users to fully realize many of the awesome features of server virtualization, such as vMotion, snapshot backups, and push button disaster recovery (to name just a few). Indeed, the most common reason for virtualizing the network is precisely to get VM mobility and vMotion to work.

In Chapter 9, “Multitenancy and the Problems of Communal Living,” you were introduced to VXLAN, which is VLAN technology with some extensions that allow it to tunnel Layer 2 frames through the IP transport network, as well as extend the number of VLANs beyond 4096. The “tunnel” that this creates allows it to bridge virtual extensible VXLAN tunnel endpoint (VTEP) devices across a network, making data transfers easy and simple regardless of where the endpoints reside (or if they move).

As noted earlier, the VMs supporting applications or services require network connectivity via physical switching and routing to be able to connect to other VMs switching the data center or cloud and with clients of the data center over a WAN link or the Internet. In addition, in a data center environment, the network also requires security and load balancing. The first switch encountered by traffic leaving the VM is the virtual switch (hypervisor), and from there a physical switch that is either top of rack (TOR) or end of row (EOR). In other words, once traffic leaves the hypervisor, it is on the physical network, and unfortunately that network cannot easily keep up with the rapidly shifting state of the VMs that are connected to it.

The way around this issue is to create a logical network of VMs that spans the physical networks the traffic travels across. VXLAN (see Figure 10-2) does this just as most network virtualization does—through the use of encapsulation. Unlike simple VLANs, though, which are limited to 4096 of these logical networks on any given physical network, VXLAN can create about 16 million. That scale is important when it comes to large data centers and clouds.




Figure 10-2 VXLAN allows for millions of logical partitions across a physical network.

Imagine that you have two VM clusters on a network, and imagine that a router separates those clusters because they are in different data centers. Both clusters in this case are on different VLANs. For these two VMs to talk to each other, the traffic between them must be routed. Now suppose you want these clusters to be on the same VLAN.

As shown in Figure 10-3, by using VXLAN, you can set up a VTEP that encapsulates or wraps the VM traffic on one end for transport over the routed network and then decapsulates (strips off the wrapper) on the other end. This effectively creates a logical network between the two clusters, which now appear to be on the same switched segment of a local network.

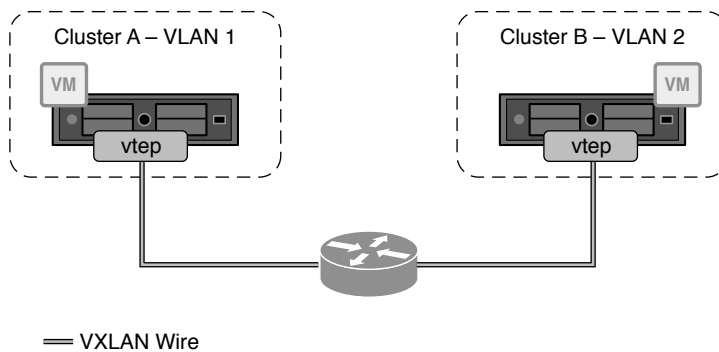


Figure 10-3 VTEPs create a logical network between the two clusters, which then appear to be on the same switched segment of a local network.

So, what's the big deal?

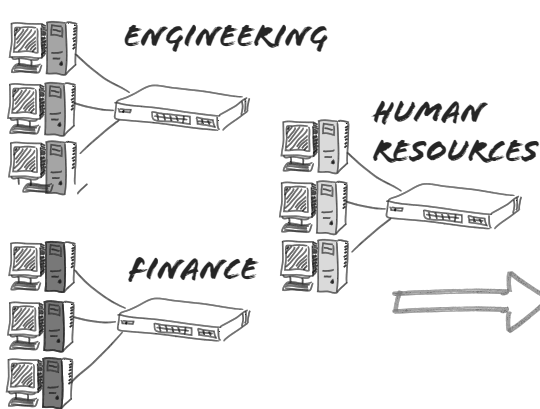
If you are new to networking, this might not seem like such a big breakthrough. If you are familiar with networking, though, you might be thinking, "This is just another way to create VLANs." There is more to it than that, though, because network virtualization in general, and VXLAN in particular, has some key benefits that become important at data center/cloud scale:

- First, this ability enables migration to a software-defined data center model. Using a vSphere administrator to provision VMs that can communicate with each other over different networks without having to involve the network team to configure the physical switches and routers eliminates one of the biggest chokepoints in the flexibility that data center virtualization affords us.
- This technology smashes through the previous limitation of 4096 VLANs
- VXLAN runs over standard switching hardware, and requires no need for software upgrades or special code versions on the switches. Therefore, you can virtualize your network using the stuff you already have.


In summary, network virtualization, although an older technology, plays a key role in the creation of virtualized data centers and clouds. It is also one of the key drivers that allows and enhances both NFV and SDN, as you will see in later chapters.

How VLANs Work

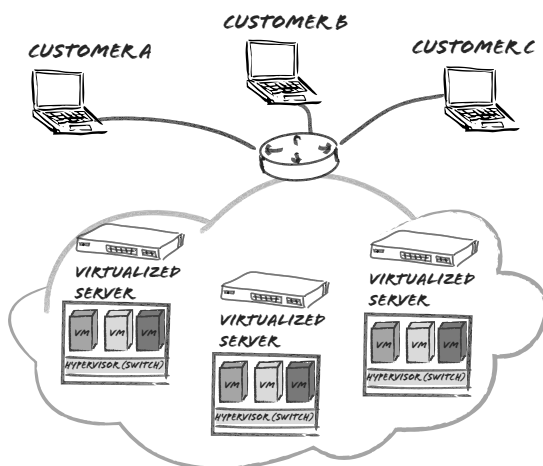
And Why They Matter in Cloud Networks



When networks were first rolled out, admins tried to keep like-users on the same switch. This local area network (LAN) model was desired, because the machines on any given LAN all received the same network instructions. It also nicely partitioned user traffic by providing physical segmentation.



As networks got bigger, it became impossible to keep like-users on a common switch. This led to the creation of the virtual LAN or VLAN. A VLAN is a means of logically grouping devices so that they appear to be connected to the same switch. VLANs provide logical partitioning and segmentation when it's not practical to physically segment.



In cloud environments, VLANs are essential because a user's VMs can span multiple servers. VLANs allow easy grouping of users resources into a single logical partition. Going further, the network admins can associate a VLAN with its IP subnet. A subnet is a way to partition a layer 3 (routing) domain. This makes it easy to connect to a cloud (and all your resources) over the Internet without needing to manually configure all of your connections.

Note that VLAN segmentation is often positioned as a form of security—VLANs do segment user data, but this is a very weak form of security that does not actually protect data.

INDEX

A

access layer, 54

access points, thin, 113

addressing

in data center, 54-55

on VMs, 55-57

VMs, MAC addresses, 56-57

adopting

NFV, 279-280

SDNs, 202-204, 277-280

advantages of multitenancy, 75

aggregation and regeneration TAPs, 257

aggregation layer, 54

aggregation, SDN use case, 189-190

agile methodology, 182

APIs, 113

northbound, 178

ODL project, 178

OnePK, 125

southbound, 178

vNetwork APIs, 266

appliances

consolidating, 210

hyper-convergence, 203-204

performance, 211

replacing with virtual appliances, 210

virtualization, 92-94

application calls, 53

Application Centric Policy, 126-128

application development

agile methodology, 182

DevOps, 182-184

NMS, SNMP, 183

pragmatic approach to, 184

SDLC waterfall methodology, 182

application layer, 116, 151

applications, 6

cloud-based, 129

availability, 130

OpenStack, 130-133

scalability, 130

networking, decoupling from gear, 147-148

one-server/application model, 13

on VMs, 129

open-source, 143

orchestration, 275

OS to application ratio, 4

provisioning, 279

recycling, 245

user-developed, 143-144

architecture

of cloud computing, 63

of data centers, 63

of ODL project, 177-178

ARP (Address Resolution Protocol), 250

ASICs, 121

attributes of cloud computing, 42

location-independent, 44

on-demand self-service, 43

pay per use, 43

rapid elasticity, 43-44

auditing

certifications for cloud providers, 230-231

TPAs, 232

automation-ready encryption, 237

availability, 217-218

as benefit of VMs, 15-16

of cloud-based applications, 130

AWS storage, 218

Azure, 230

B

bad ads, 244

bare-metal hypervisors, 22

vSphere, 68

components, 68

vMotion, 69-70

VXLAN, 70-71

Xen, 25

Microsoft Hyper-V, 26

VMware ESXi, 26

benefits

of SDNs, 187-189

of virtualization, 9, 13

availability, 15-16

ease of access to applications, 17-18

flexibility, 15-16

provisioning, 16-17

reduced cost, 13-15

reduced space, 15

standardization, 17

system administration, 17

of data center virtualization

disaster recovery, 41

easier migration, 41

faster deployment, 40

faster redeploy, 40

reduced heat, 39

reduced spending, 40

best path selection, 146

best practices for encryption in the cloud,
238-239

BFD (Bidirection Forwarding Detection), 267

block virtualization, 216

broadcast storms, 251

business cases, 188

for SDNs, 187-188, 202-203

use cases, 188

**business models for cloud computing, security
issues**, 230

BYOD (Bring Your Own Device), 235

bypass TAPs, 257

C

CAPEX, 262

financial business cases, 188

reducing with NFV, 210

Ceilometer, 132

centralized controllers for SDNs, 150-151, 169

commercial controllers, 170

NV support, 171-173

open source, 171

change control, 182

channel DLP, 225

choosing hypervisors, 27

CHS (cylinders, heads, and sectors), 215

CIA triad, 217-218, 230

Cinder, 131

**Cisco ACI (application-centric infrastructure)
Fabric**, 126-128

**Cisco APIC (Application Policy Infrastructure
Controller)**, 170

Cisco Nexus 1000V switch, 183

Cisco Nexus v1000, 266

Citrix Netscaler, 183

client-side encryption, 237

cloud computing, 41

applications, 129-130

architecture, 63

attributes of, 42

location-independent, 44

on-demand self-service, 43

pay per use, 43

rapid elasticity, 43-44

ubiquitous network access, 43

AWS storage, 218

consumers, 109

deployment models, 48

encryption

best practices, 238-239

data at rest, 236-237

data in motion, 235-236

key management, 238

Exodus Communications, 159

IaaS, 45-46, 76-78, 230

losers, 163-164

OpenStack, 130-133

overlay networks, 249

PaaS, 47, 230

providers, 109

roles, 109-110

SaaS, 44-45, 108-109, 230

multitenancy, 73-76

thin clients, 44

- security
 - audit certifications, 230-231*
 - CIA triad, 230*
 - encryption, 244*
 - logging, 231-232*
 - service creators, 109
 - versus virtualization, 129
 - winners
 - large enterprises, 160-162*
 - the “little guy”, 157-159*
 - cloud computing model, 242-243**
 - co-location, 137-138**
 - commercial controllers, 170**
 - comparing**
 - cloud computing and virtualization, 129
 - comparative and imperative SDN models, 128
 - mainframe and cloud computing models, 243
 - NFV and server virtualization, 211
 - OpenFlow and MPLS, 116
 - performance and scalability, 99-100
 - SDN and NFV, 273-274
 - VMs and NFVs, 206
 - complexity of networking, 145-146**
 - Dijkstra’s algorithm, 147
 - OSPF example, 146-147
 - components**
 - of OpenStack, 131-133
 - of vSphere, 68
 - vMotion, 69-70*
 - VXLAN, 70*
 - VTEP, 71*
 - computer models**
 - cloud computing model, 242
 - grid computing model, 242
 - Internet model, 242
 - mainframe model, 241
 - networked model, 242
 - personal computer model, 241
 - confidentiality, 217-218**
 - connectivity, data center virtualization, 53**
 - consolidated backup, 69**
 - consolidating appliances, 210**
 - consumers**
 - of cloud computing, 109
 - of virtual resources, 33
 - containers, 30**
 - control layer, 116, 152**
 - control planes, 145, 150**
 - controllers**
 - for SDNs, 169-170
 - commercial controllers, 170*
 - NV support, 171-173*
 - open source, 171*
 - ODL project, 178-179
 - OpenFlow, 116
 - conventional networks, migration to SDNs, 271-272**
 - Core Network Virtualization, 96-97**
 - cost model for data centers, 140-141**
 - COTS (commercial off the shelf) servers, 206**
 - CPE (customer premises equipment), 278**
 - CPUs, 3, 6**
 - creating a visibility layer, 267-268**
 - CRM (customer resource management), 73**
 - crypto-ransomware, 244**
 - customizing service chaining, 207**
-
- ## D
-
- data at rest, 236-237**
 - data centers**
 - architecture, 63
 - benefits of virtualizing
 - disaster recovery, 41*
 - easier migration to, 41*
 - faster deployment, 40*
 - faster redeploy, 40*
 - reduced heat buildup, 39*
 - reduced spending, 40*
 - separation of, 41*
 - server, 41*
 - design goals
 - high availability, 139*
 - low latency, 139*

- scalability, 139*
- security, 139-140*
- evolution of, 137, 142
 - co-location, 137-138*
 - cost model, 140-141*
 - design goals, 139-140*
- HDCC, 121
- networking, 53
 - addressing, 54-55*
 - application calls, 53*
 - three-tiered model, 54*
- optimization, 189, 279
- SDDC, 121
- servers, maintenance, 198-199
- switching, evolution of, 61-63
- truck rolls, scale of, 195, 198
- virtualization, connectivity, 53
- VMs
 - addressing, 55-57*
 - MAC addresses, 56-57*
- data in motion, 235**
- data leakage, preventing, 223-224**
- data loss, minimizing, 224-225**
- data sovereignty, 215**
- declarative model of forwarding control, 128**
- decoupling networking applications from gear, 147-148**
- deployment models for cloud computing, 48**
- described data, 224**
- design goals of data centers**
 - high availability, 139
 - low latency, 139
 - scalability, 139
 - security, 139-140
- DevOps (development and operations), 16, 182-184**
- DHCP (Dynamic Host Configuration Protocol), 56**
- direct inter-connects, 279**
- disaster recovery, 16, 41, 137**
- disruptive technologies, economic value of innovation, 164**
- distributed resource scheduler, 69**
- distributed virtual routing, 266**

- DLP (data leak protection), 256**
 - channel DLP, 225
 - DPI, 225, 263
 - full suite DLP, 225
 - implementing, 226
- downtime, maintenance windows, 198**
- DPI (deep packet inspection), 225, 263**
- Dropbox, 218**
- dynamic interconnects, SDN use case, 189**

E

- EaaS, 237**
- ease of access to applications, as benefit of VMs, 17-18**
- economic value of increased innovation, 164**
- edge virtual bridging, 266**
- EIGRP (Enhanced Interior Gateway Routing Protocol), 170**
- elasticity as attribute of cloud computing, 43-44**
- encapsulation, 252**
 - visibility, 274
 - VxLAN, 267
- encryption, 229, 244**
 - automation-ready, 237
 - best practices, 238-239
 - client-side, 237
 - data at rest, 236-237
 - data in motion, 235
 - EaaS, 237
 - key management, 238
 - outsourcing, 237
 - SSL, 235
 - VPN tunnels, 235
- EOR switches, 63**
- establishing SOD, silos, 182**
- ESXi hypervisor, 68**
- ETSI (European Telecommunications Standards Institute), 283**
- evolution of data centers, 137, 141-142**
 - co-location, 137-138*
 - cost model, 140-141*
 - design goals, 139-140*
 - switching, 61-63*
- Exodus Communications, 159**

F

Facebook messaging platform, 245
 fault tolerance, 69
 features of SDN, 255
 financial business cases, 188
 firewalls, 90
 flexibility, 9, 15-16
 Floodlight controller, 171
 Flowvisor controller, 171
 flow tables, 115
 flows, 170
 forwarding planes, 145
 forwarding rules manager (ODL), 179
 full suite DLP, 225
 functions of hypervisors, 24

G

Glance, 132
 goals of OpenFlow, 116
 green screens, 241
 GRE (generic routing encryption) tunnels, 249
 grid computing model, 242

H

hairpinning, 122
 hardware resources of servers, 6-7
 HDDC (hardware-defined data center), 121
 header rewrite, 172
 headers, 144
 Heat, 132
 heat generation of servers, 5
 Helium, 175
 hierarchical network model, 53
 hierarchy of orchestrators, 275
 high availability as data center design goal, 139
 HIPAA (Health Insurance Portability and
 Accountability Act), 230
 history
 of hypervisors, 22
 of OpenFlow, 114-115

Horizon, 131
 hosted hypervisors, 23
 host tracker (ODL), 179
 Hotmail, 159
 HP VAN (Virtual Application Networks)
 Controller, 170
 hybrid clouds, 48
 hybrid SDN solution, 126
 Hydrogen, 175
 hyper-convergence, 203-204
 Hyper-V Virtual Switch, 265
 hypervisor vSwitch, 65
 hypervisors, 8-9, 21
 bare-metal, 22, 25-26
 choosing, 27
 functions of, 24
 history of, 22
 hosted, 23
 monitoring function, 22
 vendors, 24
 virtual resources, managing, 31-32
 VMware, 67
 vSphere, 68
 components, 68
 vMotion, 69-70
 VXLAN, 70-71

I-J

IaaS (Infrastructure-as-a-Service), 45-46, 76-78,
 230, 242
 IBM PC, 241
 IDS (intrusion detection systems), 256
 imperative model of forwarding control, 128
 implementing DLP, 226
 importance of SDN, 285
 in-band signaling, 268
 infrastructure layer, 116, 152
 innovation
 economic value of, 164-165
 user-developed applications, 143
 Insieme, 126-128

integrating

- physical and overlay networks, 267–268
- SDN and NFV, 273–274

integrity, 217–218**interest in NFV, reasons for, 283****internal fight for network control, 181****Internet model, 242****inter VM-to-VM traffic, monitoring, 266****intra-VM traffic, monitoring, 265****IP addressing, VMs, 56****IPS (intrusion prevention systems), 256**

K**key management, 238****Keystone, 131****KPIs (key performance indicators), 173****KVM, 24**

L**Layer 2 networks, scalability, 250–251****layer 4 through 7 network services, 89**

- firewalls, 90
- load balancers, 90–92
- SSL offload, 90
- VPNs, 90

layout

- of cloud computing, 63
- of data centers, 63

LBA (logical block addressing), 215**link state, 146****Lithium, 175****load balancers, 90–92****location-independent resource pooling, 44****log collectors, 232****logging, 231–232****logical view of SDNs, 150****loops in Layer 2 networks, 251****low latency as data center design goal, 139****LSPs (label switch paths), 206****LTE (Long-Term Evolution), 95**

M**MAC addresses for VMs, 56–57****mainframe model, 241**

- comparing to cloud computing model, 243
- security, 244

maintenance of servers, 198–199**maintenance windows, 198****malware, 244****managing virtual resources, 29–34****manifesto, ODL project, 176–177****Meerkat, 245****memory, 6****merging application development and operations, 182****Microsoft Hyper-V, 26****migrating to SDN, 202–204**

- hyper-convergence, 203–204
- reasons for, 271–272

minimizing data loss, 224–225**mobile apps, recycling, 245****mobile service providers, Core Network Virtualization, 96–97****mobility virtualization, 278****monitoring function of hypervisors, 22****monitoring networks**

- inter VM-to-VM traffic, 266
- intra-VM traffic, 265
- SPAN ports, 256
- TAPs, 256–257
- traffic, SDN's lack of network management, 274

monitoring tools, 256**Moore's Law, 65****MPLS (Multi-Protocol Label Switching), 83, 116, 249–250**

- LSPs, 206
- versus OpenFlow, 116

multiplexing, 22**multitenancy, 44, 73–74**

- advantages of, 75
- IaaS, 76–78
- security risks, 76
- shared multitenant clouds, 48

N

NAC, 189, 279

NAT (Network Address Translation), 114

NEC ProgrammableFlow PF6800 Controller, 170

network analyzers

SPAN ports, 256

TAPs, 256–257

network breakout TAPs, 257

network management tools

SPAN ports, 256

TAPs, 256–257

network optimization, 278

network services, layer 4 through 7, 89–92

firewalls, 90

load balancers, 90–91

SSL offload, 90

VPNs, 90

network virtualization, 272, 279

network visibility, 255

networked model, 242

networking, 144–145

applications, decoupling from gear, 147–148

best path selection, 146

complexity, 145–146

Dijkstra's algorithm, 147

OSPF example, 146–147

control planes, 145

Dijkstra's algorithm, 147

in traditional data centers

addressing, 54–55

application calls, 53

three-tiered model, 54

OSI stack, 144

packets, 144

payload, 144

routers, 144

routing protocols, OSPF, 146–147

switches, 144

networking forwarding planes, 145

networks

monitoring, network speed effect on, 256

SDN, 85

virtualization, 83–87

performance, 100–102

scalability, 101–102

Neutron, 131

Nexus v1000 vSwitch, 266

NFV (network functions virtualization), 85, 209

adopting, 279–280

and SDN, 283

appliances, 210

as CAPEX-reducing technology, 210

business suitability for, 201

deploying, 277

ETSI, 283

integration with SDN, 273–274

interest in, reasons for, 283

performance, 211

proof of concepts, vCPE, 284

role in complementing SDN, 285

underlay networks, 273

versus server virtualization, 211

versus VMs, 206

winners

large enterprises, 160–162

the “little guy”, 157–159

Nicira, 119

VMware NSX, 120

NV, 120–122

promiscuity, 123–124

NIC (network interface card), 6

NIST (National Institute for Science and Technology), 42

NMS (network management software), 183

inter-VM-to-VM traffic, monitoring, 266

intra-VM traffic, monitoring, 265

orchestration, 274–275

SNMP, 183

visibility, 184, 274

NOC (network control center), 183

northbound APIs, 178

Nova, 131

NPBs (network packet brokers), 258

NSX platform, 120

NV, 120–122

promiscuity, 123–124

NV (network virtualization), 93-95, 120-124

controller support for, 171-173

SDN use case, 189

VMs, 138

NVGRE, 252

O**ODL (OpenDaylight) project, 175**

architecture, 177-178

controller platform, 178-179

Lithium, 175

manifesto, 176-177

releases, 175

on-demand self-service, as attribute of cloud computing, 43**OnePK, 125****one-server/application model, 13**

benefits of

*availability, 15-16**ease of access to, 17-18**flexibility, 15-16**provisioning, 16-17**reduced cost, 13-15**reduced space, 15**standardization, 17**system, 17***ONF (Open Network Foundation), 175****Open Networking Foundation, 117****open source, 130**

applications, 143

ODL project, 175

*architecture, 177-178**controller platform, 178-179**Lithium, 175**manifesto, 176-177**releases, 175*

ONE, 175

OpenFlow, 113-114, 125

controllers, 116

flow tables, 115

goal of, 116

history of, 114-115

versus MPLS, 116

OpenContrail SDN controller, 171**OpenDaylight open source SDN controller, 171****OpenStack, 119, 129-133****operational domains**

mobility virtualization, 278

network optimization, 278

provisioning, 279

service orchestration, 278

virtual CPE, 278

WAN optimization, 278

OPEX, 262**OpFlex, 126****orchestration, 274-275****OS (operating systems), 6-9****OS-to-application ratio, 4****OSI stack, 144****OSPF, 146-147****OSS/BSS (operations systems support/business support system), 284****out-of-band signaling, 268****outsourcing encryption, 237****overlay networks, 249**

encapsulation, 252

MPLS, 249-250

integrating with physical network, 267-268

tunnels

*provisioning with SDN, 252***overprovisioning, 256**

P**PaaS (Platform-as-a-Service), 47, 230, 242****packets, 144****para-virtualization, 25****pay per use, as attribute of cloud computing, 43****payload, 144****PBB (provider backbone bridging), 172****PCI-DSS, 230****performance**

in network virtualization, 100-102

of appliances, 211

orchestration, 274-275

versus scalability, 99-100

Periscope, 245

personal computer model, 241

physical network, integrating with overlay network, 267-268

Pluggable Storage Array, 69

preventing

- data leakage, 223-224
- data loss
 - channel DLP*, 225
 - DPI*, 225
 - full suite DLP*, 225

private clouds, 48, 243

programmability

- customizing service chaining, 207
- of SDNs, 153-154, 285

proof of concepts for NFV, vCPE, 284

providers

- of cloud computing, 109
- of virtual resources, 33
- virtual customer edge, SDN use case, 189

provisioning, 279

- as benefit of VMs, 16-17
- overprovisioning, 256
- self-service, 243
- tunnels, 252
- VLANs, 272

public clouds, 48

Q-R

QoE (quality of experience), 259-263

QoS (quality of service), 114, 172, 255

RAM, 6

ransomware, 244

recycling apps, 245

registered data, 224

regulations

- HIPAA, 230
- PCI-DSS, 230
- security, 230
- SOX, 172

releases of ODL project, 175

replacing appliances with virtual appliances, 210

ROI (return on investment), 31

roles in cloud computing, 109-110

routers, 144

routing protocols

- Dijkstra's algorithm, 147
- OSPF, 146-147

Ryu OpenFlow Controller, 171

S

SA 300 standard, 230

SA 315 standard, 230

SaaS (Software-as-a-Service), 45, 108-109, 230, 242

- multitenancy, 73-74
 - cost savings*, 75
 - security risks*, 76
- thin clients, 44

sandboxing, 14

SANs (storage-area networks), 17, 216-217

scalability

- as data center design goal, 139
- of cloud-based applications, 130
- of Layer 2 networks, 250
 - broadcast storms*, 251
 - loops*, 251
- NV, 101-102
- versus performance, 99-100

scale of truck rolls, 195, 198

SDDC (software-defined data center), 121

SDLC (software development lifecycle), 182

SDNs (software-defined networks), 10, 85, 113, 149, 235

- adopting, 202-203, 279-280
- and NFV, 283
- application layer, 151
- benefits of, 187-189
- business case for, 187-188
- business suitability for, 201-202
- centralized controller, 150-151
- Cisco ACI, 127-128
- control layer, 152
- controllers, 169-173
- control plane, 150

- deploying, 277
- early implementations of, 185
- features, 255
- hybrid SDN, 126
- importance of, 285
- infrastructure layer, 152
- Insieme, 126
- integration with NFV, 273-274
- lack of network management visibility, 274
- logical view, 150
- migration to, reasons for, 271-272
- network virtualization, 272
- Nicira, 119
- OnePK, 125
- ODL project, 175
 - architecture, 177-178*
 - controller platform, 178-179*
 - Lithium, 175*
 - manifesto, 176-177*
 - releases, 175*
- OpenFlow, 114
 - controllers, 116*
 - flow tables, 115*
 - goal of, 116*
 - history of, 114-115*
 - versus MPLS, 116*
- orchestration, 274-275
- programmability, 153-154
- role in complementing NFV, 285
- service chaining, 205-207
- traffic flows, 150, 170
- tunnels, provisioning, 252
- use cases
 - data center optimization, 189*
 - DCO, 279*
 - dynamic interconnects, 189*
 - NAC, 189*
 - NAC, 279*
 - NV, 189*
 - NV, 279*
 - virtual core and aggregation, 189-190*
 - virtual customer edge, 189*
- virtualization layer, 152-153
- visibility, need for, 255
- VLANs, provisioning, 272
- winners
 - large enterprises, 160-162*
 - the "little guy," 157-159*
- security, 244**
 - as data center design goal, 139-140
 - bad ads, 244
 - CIA triad, 217-218, 230
 - data leakage, preventing, 223-224
 - data sovereignty, 215
 - DevOps, 184
 - encryption, 229, 235, 244
 - automation-ready, 237*
 - best practices, 238-239*
 - data at rest, 236-237*
 - data in motion, 235*
 - key management, 238*
 - outsourcing, 237*
 - in cloud computing
 - audit certifications, 230-231*
 - logging, 231-232*
 - multitenancy, 76-78
 - ransomware, 244
 - regulations, 230
 - service chaining, 205-207
- segmentation, 114**
- selecting hypervisors, 27**
- self-service, 243**
- server virtualization versus NFV, 211**
- servers**
 - applications, 6
 - hardware resources, 6-7
 - heat generation, 5
 - maintenance, 198-199
 - OS, 6-9
 - proliferation of, 4-5
 - resources, 6
 - truck rolls, 195
 - underutilization, 7
 - virtualization, 85

service chaining, 205-207
service creators of cloud computing, 109
service orchestration, 278
service providers
 operational domains
 mobility, 278
 network optimization, 278
 provisioning, 279
 service, 278
 virtual CPE, 278
 WAN optimization, 278
 QoE, 261-263
 SDN adoption, 277
SFDC (SalesForce.com), 73-74
shared multitenant clouds, 48
SIEM (security and event management system), 231
silos, 182
SLAs (service-level agreements), 236, 274-275
SNMP (Simple Network Management Protocol), 183
SOA (service-oriented application), 280
SOD (separation of duties), 181-182
southbound APIs, 178
SOX (Sarbanes-Oxley), 172, 230
SPAN (Switched Port Analyzer) ports, 256
SPB (shortest path bridging), 172
SRS (software requirement specification), 182
SSAE 16 standard, 230
SSL encryption, 235
SSL offload, 90
SST (Single Spanning Tree), 252
standardization as benefit of VMs, 17
statistic manager (ODL), 179
storage, 6
storage virtualization
 AWS storage, 218
 block virtualization, 216
 LBA, 215
Storage vMotion, 70
STP (Spanning Tree Protocol), 71, 251
suitability
 for NFV, 201
 for SDN, 201-203

Swift, 131
switches, 144
 data centers, evolution of, 61-63
 EOR, 63
 hypervisor vSwitch, 65
 SPAN ports, 256
 TAPs, 256-257
 TOR, 63
 virtualized aware network switches, 65
SysOps (systems and operations), 16
system administration as benefit of VMs, 17

T

TAPs (test access points), 256-257
thin access points, 113
thin clients, 44
three-tiered networking model, 54
TLF (The Linux Foundation), ODL project, 175
 architecture, 177-178
 controller platform, 178-179
 Lithium, 175
 manifesto, 176-177
 releases, 175
tool farms, 258
topology manager (ODL), 179
TOR switches, 63
TPAs (third-party auditors), 232
traditional networks, migration to SDNs, 271-272
traffic flows
 in SDNs, 150
 ODL, 177
TRILL (transparent interconnection of lots of links), 172
Trove, 132
truck rolls, 195, 198
TTL (time-to-live) counter, 251
tunneling
 MPLS, 249-250
 overlay networks, 249
 provisioning with SDN, 252
 VEP, 253
 visibility, 274

VNIDs, 253

VxLAN, 267

Type 1 hypervisors, 22

Type 2 hypervisors, 23

U

ubiquitous network access as attribute of cloud computing, 43

UDP (User Datagram Protocol), VxLAN, 267

underlay networks, 249, 273

underutilization of servers, 7

use cases, 188

financial business cases, 188

for SDNs, 202–203

CDO, 279

data center optimization, 189

dynamic interconnects, 189

NAC, 189, 279

NV, 189, 279

virtual core and aggregation, 189–190

virtual customer edge, 189

user-developed applications, 143–144

V

VASs (value-added services), 97

vCenter Server, 68

vCPE (virtual customer premises equipment), 278, 284

VDI (virtual desktop infrastructure), 17

VDS (Virtual Distribution Switch), 265

vendors of hypervisors, 24

VEPA (Virtual Ethernet Port Aggregator), 266

vEPC (evolved packet core), 189

VEPs (virtual end points), 253

vIMS (virtual Information Management System), 189

virtual customer edge, SDN use case, 189

virtual overlays, 258–259

virtual resources

consumers, 33

managing, 29–34

providers, 33

Virtual SMP, 69

virtualization, 3

appliances, 92–94

benefits of, 9

availability, 15–16

ease of access to applications, 17–18

flexibility, 15–16

provisioning, 16–17

reduced cost, 13–15

reduced space, 15

standardization, 17

system administration, 17

block virtualization, 216

Core Network Virtualization, 96–97

early implementations of, 185

hypervisors, 8–9, 21

bare-metal, 22

choosing, 27

functions of, 24

history of, 22

hosted, 23

monitoring function, 22

vendors, 24

Xen, 25

networks, 83–87

performance, 100–102

VXLAN, 100

NFV, 85

NV, 95, 189, 272

servers, 85, 195

server virtualization, versus NFV, 211

storage virtualization, 215

versus cloud computing, 129

VMs, 3, 7–10, 129

workload, 30–31

virtualization layer (SDNs), 152–153

virtualizing the data center, benefits of

disaster recovery, 41

easier migration, 41

faster deployment, 40

faster redeploy, 40

reduced heat, 39

- reduced spending, 40
 - visibility, 255**
 - creating, 267-268
 - intra VM-to-VM traffic flows, 265
 - in tunnels, 274
 - NMS, 184
 - with SDN, 274
 - visibility layer, creating, 257**
 - VLANS, 83, 249, 272**
 - VM-aware switching, 65**
 - VMFS (Virtual Machine File System), 69**
 - vMotion, 13-15, 69-70**
 - VMs (virtual machines), 3, 7-10, 21, 138**
 - addressing, 55-57
 - applications, 129
 - benefits of
 - availability, 15-16*
 - ease of access to applications, 17-18*
 - flexibility, 15-16*
 - provisioning, 16-17*
 - reduced cost, 13-15*
 - reduced space, 15*
 - standardization, 17*
 - system administration, 17*
 - inter VM-to-VM traffic, monitoring, 266
 - intra-VM traffic, monitoring, 265
 - IP addressing, 56
 - MAC addresses, 56-57
 - versus NFVs, 206
 - visibility, 265
 - visibility layer, creating, 267-268
 - VMware, 67
 - zombies, 184
 - VMware, 67**
 - ESXi, 26
 - NSX platform, 120-122
 - NV, 120-121*
 - promiscuity, 123-124*
 - vSphere, 68
 - components, 68*
 - vMotion, 69-70*
 - VXLAN, 70-71*
 - VMware Server, 68**
 - vNetwork APIs, 266**
 - vNetwork Distribution Switch, 69**
 - vNICs, 56, 265**
 - VNIDs (virtual network IDs), 253**
 - VPNs (virtual private networks), 90, 235**
 - VRF (Virtual Routing and Forwarding), 171**
 - vSphere, 68, 183**
 - vMotion, 69-70*
 - VXLAN, 70-71*
 - vSwitches**
 - Cisco Nexus 1000V switch, 183, 266
 - distributed virtual switching, 266
 - VTEPs (virtual tunnel endpoints), 71, 258**
 - VXLAN (Virtual Extensible LAN), 70-71, 86, 100, 252, 267**
-
- ## W
-
- WAN optimization, 278**
 - web apps, recycling, 245**
 - websites, OpenStack, 132**
 - Wi-Fi, “thin” access points, 113**
 - workload, 30-31**
-
- ## X-Y-Z
-
- Xen hypervisors, 25-26**
 - Zaqr, 132**
 - zombies, 184**