

LEARNING MADE EASY

Forcepoint Custom Edition

# Secure Enterprise SD-WAN

for  
**dummies**<sup>®</sup>  
A Wiley Brand



Understand  
SD-WAN pros and cons

Learn how to get  
started with SD-WAN

See what makes SD-WAN  
enterprise-ready

Brought to  
you by

 **FORCEPOINT**

Joe Kraynak

## **About Forcepoint**

Forcepoint is transforming cybersecurity by focusing on what matters most: understanding people's intent as they interact with critical data and intellectual property wherever it resides. Our uncompromising systems enable companies to empower employees with unobstructed access to confidential data while protecting intellectual property and simplifying compliance. Based in Austin, Texas, Forcepoint supports more than 20,000 organizations worldwide. For more about Forcepoint, visit [www.Forcepoint.com](http://www.Forcepoint.com) and follow us on Twitter @ForcepointSec.



# Secure Enterprise SD-WAN

Forcepoint Custom Edition

**by Joe Kraynak**

**for  
dummies®**  
A Wiley Brand

# Secure Enterprise SD-WAN For Dummies®, Forcepoint Custom Edition

Published by  
**John Wiley & Sons, Inc.**  
111 River St.  
Hoboken, NJ 07030-5774  
[www.wiley.com](http://www.wiley.com)

Copyright © 2018 by John Wiley & Sons, Inc., Hoboken, New Jersey

No part of this publication may be reproduced, stored in a retrieval system or transmitted in any form or by any means, electronic, mechanical, photocopying, recording, scanning or otherwise, except as permitted under Sections 107 or 108 of the 1976 United States Copyright Act, without the prior written permission of the Publisher. Requests to the Publisher for permission should be addressed to the Permissions Department, John Wiley & Sons, Inc., 111 River Street, Hoboken, NJ 07030, (201) 748-6011, fax (201) 748-6008, or online at <http://www.wiley.com/go/permissions>.

**Trademarks:** Wiley, For Dummies, the Dummies Man logo, The Dummies Way, Dummies.com, Making Everything Easier, and related trade dress are trademarks or registered trademarks of John Wiley & Sons, Inc. and/or its affiliates in the United States and other countries, and may not be used without written permission. All other trademarks are the property of their respective owners. John Wiley & Sons, Inc., is not associated with any product or vendor mentioned in this book.

LIMIT OF LIABILITY/DISCLAIMER OF WARRANTY: THE PUBLISHER AND THE AUTHOR MAKE NO REPRESENTATIONS OR WARRANTIES WITH RESPECT TO THE ACCURACY OR COMPLETENESS OF THE CONTENTS OF THIS WORK AND SPECIFICALLY DISCLAIM ALL WARRANTIES, INCLUDING WITHOUT LIMITATION WARRANTIES OF FITNESS FOR A PARTICULAR PURPOSE. NO WARRANTY MAY BE CREATED OR EXTENDED BY SALES OR PROMOTIONAL MATERIALS. THE ADVICE AND STRATEGIES CONTAINED HEREIN MAY NOT BE SUITABLE FOR EVERY SITUATION. THIS WORK IS SOLD WITH THE UNDERSTANDING THAT THE PUBLISHER IS NOT ENGAGED IN RENDERING LEGAL, ACCOUNTING, OR OTHER PROFESSIONAL SERVICES. IF PROFESSIONAL ASSISTANCE IS REQUIRED, THE SERVICES OF A COMPETENT PROFESSIONAL PERSON SHOULD BE SOUGHT. NEITHER THE PUBLISHER NOR THE AUTHOR SHALL BE LIABLE FOR DAMAGES ARISING HEREFROM. THE FACT THAT AN ORGANIZATION OR WEBSITE IS REFERRED TO IN THIS WORK AS A CITATION AND/OR A POTENTIAL SOURCE OF FURTHER INFORMATION DOES NOT MEAN THAT THE AUTHOR OR THE PUBLISHER ENDORSES THE INFORMATION THE ORGANIZATION OR WEBSITE MAY PROVIDE OR RECOMMENDATIONS IT MAY MAKE. FURTHER, READERS SHOULD BE AWARE THAT INTERNET WEBSITES LISTED IN THIS WORK MAY HAVE CHANGED OR DISAPPEARED BETWEEN WHEN THIS WORK WAS WRITTEN AND WHEN IT IS READ.

For general information on our other products and services, or how to create a custom *For Dummies* book for your business or organization, please contact our Business Development Department in the U.S. at 877-409-4177, contact [info@dummies.biz](mailto:info@dummies.biz), or visit [www.wiley.com/go/custompub](http://www.wiley.com/go/custompub). For information about licensing the *For Dummies* brand for products or services, contact [BrandedRights&Licenses@Wiley.com](mailto:BrandedRights&Licenses@Wiley.com).

ISBN: 978-1-119-50927-1 (pbk); ISBN: 978-1-119-50926-4 (ebk)

Manufactured in the United States of America

10 9 8 7 6 5 4 3 2 1

## Publisher's Acknowledgments

### Development Editor:

Christina N. Guthrie

### Project Editor:

Christina N. Guthrie

### Acquisitions Editors:

Amy Fandrei,  
Steven Hayes

### Editorial Manager:

Rev Mengle

### Business Development

Representative: Molly Daughtery

### Production Editor:

Vasanth Koilraj

# Introduction

It's the 21st century. Networking should work, hassle-free and cost-effectively. You should be able to connect all your geographically dispersed locations and computing resources, including cloud-based resources, seamlessly without having to worry about the privacy of data transmissions. You should be able to aggregate multiple connections of various types to create reliable high-speed connections affordably. And you should be able to monitor and control your entire wide area network (WAN) from a central location through a single pane of glass.

*Software-defined wide area networking (SD-WAN)* delivers all of these capabilities. The result is a fast, reliable, and secure WAN that's easy to deploy and manage. Secure Enterprise SD-WAN takes this even further, providing an integrated way to both connect and protect locations across your organizations more efficiently and effectively than ever before.

## About This Book

In this book, we bring you up to speed on networking and WAN and the challenges you're likely to encounter. This enables you to communicate more effectively with your technology personnel and vendors. We then set the context for SD-WAN by examining the evolution of information technology (IT). Finally, we explain what SD-WAN is and how it works and provide guidance on how to bring your organization into the 21st century with the latest in Secure Enterprise SD-WAN technology.

## Foolish Assumptions

While writing this book, we made some assumptions about you:

- » You're part of a large or mid-sized organization where some of or all the activity is digital.
- » You're familiar with IT and have some knowledge of how operations are managed in your organization.
- » You're interested in understanding what different options you can use to manage a growing network.

» You have a proactive approach to IT and want to discover how to keep abreast of changing and disparate technologies.

## Icons Used in This Book

We use the following icons to highlight key text so that you can navigate easily to the most useful information:



CASE STUDY

This icon points out real-world examples you may find helpful.



REMEMBER

Here we highlight important information for you to bear in mind.



TIP

This icon draws your attention to top-notch advice.



WARNING

Watch out for these potential pitfalls!

## Beyond the Book

In a short book like this there is only so much we can cover. So, if you find yourself wanting to know more about hybrid IT, and Secure Enterprise SD-WAN in particular, just go to [www.forcepoint.com/ngfw](http://www.forcepoint.com/ngfw).

## Where to Go from Here

You can use this book however you like. By all means take the traditional route and read it straight through from start to finish. Or you can skip between sections or chapters, using the chapter titles and section headings as your guide to pinpoint the information you need. Whichever way you read it, you can't go wrong. All paths lead to the same outcome: a better grasp of how the right technology works to make large and diversified networks more agile and more secure.

- » Getting up to speed on Internet protocols
- » Directing network traffic with switches and routers
- » Ensuring service availability
- » Grasping network security basics
- » Understanding network performance terminology

# Chapter 1

## Brushing Up on Networking Technology and Terminology

In society and in business, people need to communicate their ideas and share information in order to interact, collaborate, and create. People have used technology for some time to achieve this goal, technology that has evolved rapidly from the telephone to email to social media. Regardless of the mechanism, the process is always the same: a person who possesses information or data passes it to another person or persons. As part of doing so, they may make it public — for example, using a bullhorn to address an audience — or they may keep it private, perhaps by having a quiet conversation with no one else present. They may want to communicate with a specific person or speak with anyone from, say, the customer service department. Perhaps they wish to communicate privately with a group of people they trust. This concept of managing communication — its scope, reach, and privacy — applies equally within the world of information technology (IT) and is enabled by networking technology.

A network, in its simplest form, allows multiple devices to share data and resources. Its job is to make sure that the data being sent from one device to another gets there intact, in the shortest possible time and without interfering with any other data transfers. As with communication among people, the network must be capable of enabling communication between two or more devices, regardless of their locations, while protecting the integrity and confidentiality of that data. These challenges have driven the evolution of networking technology since its inception and the development of the first wide area network (WAN) known as ARAPANET in the 1960s.

*Software-defined wide area networking* (SD-WAN) seamlessly connects users in geographically dispersed locations to one another and to both organizational and cloud resources. With SD-WAN, businesses essentially have one big network with centralized control. SD-WAN was created to harness the capabilities of existing network technology, including commodity broadband and virtual private networks (VPNs), and extend them with better traffic management and monitoring as well as aligning their operation with business policies. (Commodity broadband's cost-effective performance can come in different forms, depending upon the Internet access services available in a given location. VPN technology employs encryption to secure connections to remote computers.) The result is a better networking solution that's faster to deploy, all at a lower cost compared to traditional approaches.

Before we dive into this exciting technology, we need to first explore the factors that are driving the need for it. In this chapter, we introduce you to the elements of networks and the overarching goals of networking in a decidedly nontechnical manner. We dig deeper into the technology underlying the magic and introduce you to the terminology you need to know to have a productive conversation about network function, availability, security, and performance. Armed with the knowledge of networking, you'll be well equipped to explore SD-WAN and how Secure Enterprise SD-WAN approaches are changing how enterprises connect and protect their organizations.



# How Network Devices Connect and Communicate

A variety of technologies and standards determine how network devices connect and communicate with one another. To understand how data travels from point A to point B across a local network or across a wide-area network (especially one that uses the Internet), it helps if you understand these technologies and standards and the terminology used to refer to them.

## Introducing TCP/IP

*Network protocols* are collections of rules and conventions for communicating and transferring data between devices on a network. You can think of them as being like the rules that govern the English language, such as grammar, usage, syntax, and punctuation, enabling people to communicate with one another. Without such rules, people would be unable to communicate effectively. Likewise, without protocols, networked devices would be unable to connect and communicate with one another.

In this section, we describe a few of the most common network protocols and groups of protocols.

## Internet protocol (IP)

*Internet Protocol* (IP for short) is a large suite of interrelated protocols that collectively govern how modern networks are created and operate. IP contains rules for the format of the data being passed, how it's addressed, how destinations are discovered and reached, and more.

Internet Protocol sets the standard for moving a *packet* between two networked devices, wherever they may be. Other protocols are used by the sending device to determine whether the sender can find the recipient directly (*switching*) or whether it will need help from a specialized device to do so (*routing*). See the later section, “Directing traffic with switches and routers” for details.

Internet protocol is available in two versions: *IPv4* and *IPv6*, each of which provides addresses for the various devices connected on a network. Key differences between the two are the address format and the number of addresses they support:

- » **IPv4** uses 32-bit numeric addresses written as four groups of numbers, using periods to separate the groups; for example, 164.8.41.4. It supports up to 4.3 billion unique addresses.
- » **IPv6** uses 128-bit hexadecimal addresses written as eight groups of four numbers, with each group containing letters, numbers, or both; for example, 2001:0DBB:AC10:FE01:3FFE:4545:FE21:67CF. IPv6 can support  $3 \times 10^{38}$  (three trillion trillion trillion) addresses, which certainly seems like overkill, but who knows how many devices will eventually be connected across the Internet? *Hexadecimal* refers to 16 digits, represented by 0–9 (ten numbers) and A–F (six letters); A = 10, B = 11, C = 12, up to F = 15.

IPv6 is actually the successor to IPv4. As more and more devices have connected to the Internet, the system actually ran out of IPv4 addresses, so IPv6 was introduced to supply more addresses and some additional benefits. However, many organizations still use IPv4 internally, so you're likely to encounter IPv4 addresses for some time to come.



REMEMBER

The good news is that you rarely need to use IP addresses, because they generally operate behind the scenes. You enter “www.wiley.com” into your Web browser, and the Internet directs your request to the IP address of the computer associated with that domain name. You address an email message to a friend or colleague, and the Internet carries that message to the IP address associated with the recipient's email server, which then places that message in the recipient's inbox. You're likely to deal with IP addresses only if you're configuring or troubleshooting your network.

## Transmission Control Protocol/Internet Protocol (TCP/IP)

TCP/IP is the formal name (usually it is just referred to as TCP) for another network protocol that runs on top of the IPv4 or IPv6 protocols mentioned above to provide a reliable stream of

communication between two devices on a network. TCP and IP each have a distinct role:

- » **TCP** sets the rules for how applications create communication channels across the network, how a message is broken into packets before being sent, and how those packets are reassembled after reaching their destination.
- » **IP** sets the rules for how each packet is addressed and routed so it reaches the right destination. As a packet passes through various network devices on the Internet, each device checks each packet's IP address, so it can forward it to the right destination.

TCP/IP is all about increasing the reliability of communication across the Internet. However, its strength of ensuring delivery somewhat limits communication performance, because the device on the receiving end must confirm receipt of the entire message. If one or more packets don't arrive, then the sending device must resend them. Degradation of performance can impact, for example, how quickly a web page loads or the quality of a videoconferencing call. Some applications can fail altogether. In addition, TCP/IP is of no use when a user's connection with the Internet service provider (ISP) goes down.



REMEMBER

Network issues that commonly impact the ability of TCP/IP to do its job include packet drop and latency. *Packet drop* (also referred to as *packet loss*) occurs when a packet fails to reach its destination, primarily due to network congestion. Network latency is an expression of how much time it takes for a packet of data to get from one designated point to another. In some environments, latency is measured by sending a packet that is returned to the sender; the round-trip time is considered the latency. Ideally, latency is as close to zero as possible.

## Directing network traffic with switches and routers

As devices exchange data across the Internet, the data bounces around from computer to computer, from one network to another, until it eventually reaches its destination. All day every day, massive volumes of data are transmitted, and all of it is broken down into packets. Someone, or something, has to be in charge of directing all this traffic. Two types of devices are responsible for directing traffic: routers and switches.

## Routing data between networks

Routers forward packets between devices in different networks over layer 3 of the network. A *layer* is a distinct network function. Layer 3 handles the addressing and routing of data through the use of IP addresses.

Routers are equipped with special software that implements standard routing protocols that automatically make decisions of where to send a packet. The routing protocol builds a map, and the router chooses the best path to take. After a router figures out how to reach a destination network, it “remembers” the destination, greatly speeding up future communication.

Routing protocols typically require configuration on the part of network administrators. Misconfiguration can cause network disruptions and has been the culprit behind some of the largest outages of public cloud services.



REMEMBER

Switches, discussed in the next section, have optional routing features but can't be used to connect to the Internet for all but the smallest networks. They're not built to learn and remember routes at the same scale as routers specifically designed for this purpose, nor do they perform routing at the same speed.

## Directing traffic with switches

Switches are network devices that forward packets between devices in the same network. The difference between a switch and a router is that a switch connects devices on a network, whereas a router connects networks. Network switches are most commonly used in business networks to connect computers, printers, servers, and other computer resources within an office or building.

## Grasping Service Availability Basics

A key challenge to managing any network is to keep it up and running so it's always available for people who need to use it. The following three tools are very helpful in overcoming this challenge:

» **Load balancing:** A load balancer is a device that takes a request received via its well-known IP address and forwards it to one of a pool of servers that has the available capacity



TIP

to respond to the request. A load balancer is a key to keeping a large website up and running and ensuring that it can quickly respond to user requests. Load balancing is sometimes provided as a feature within other network devices like firewalls.

You can also apply load balancing within a network to distribute traffic among multiple wide area network (WAN) providers, a technique we describe in greater detail in Chapter 3.

- » **Redundancy:** Redundancy involves duplicating equipment and connections to provide multiple paths between any two endpoints over different connections and different equipment. Historically, the weak link in networks, particularly branch offices, was Internet access either through the home office or through a single Internet Service Provider (ISP). If that connection went down, that branch would be disconnected from the WAN. See Chapter 3 for additional details.
- » **Clustering:** With clustering (a form of redundancy), you connect multiple computing resources in remote locations to create what functions as a single unified network with high availability.
- » **Service level agreement (SLA):** An SLA is a performance and availability guarantee. Network service providers that offer managed point-to-point connections (MPLS) typically provide SLAs to ensure customers that the network will be available a high percentage of the time and will perform as expected. By contrast, although the Internet is highly reliable, most ISPs offer customers no guarantee that they'll be able to connect at any given time or at any given speed.

## Checking Out Network Security Options

Networking security or lack thereof makes the headlines regularly with stories of evil hackers breaking into networks and stealing data, infecting networks with malware, and spying on users. It's what keeps network managers and cyber security experts awake

at night. Fortunately, you have some options for protecting your network and your data both at rest and when it's in transit. Here are a few key options:

» **Virtual local area network (VLAN):** The purpose of a LAN is to enable devices on the network to connect and communicate seamlessly across the network connections. However, in some cases, you want certain parts of your network to be more isolated and protected. For example, suppose you have your Sales and Finance departments on the same floor of an office building. You want the Sales computers to connect and communicate freely with one another and the Finance computers to connect and communicate freely with one another, but you don't want your salespeople to have unfettered access to financial data. Putting each department on its own physical network can be a hassle, if not impossible, if all their cubicles are connected to the same switch. VLANs resolve this issue by using virtual switches within the physical switch, creating the appearance of two physically separate networks. You can then tighten security on Finance networks without affecting accessibility on the Sales network.

» **Virtual private network (VPN):** A broad class of protocols, including IPsec, SSL, VPLS, EVPN, EVPL, and GRE, can create a private network connection for customers across otherwise public or shared WAN infrastructure, including the Internet. Imagine a tunnel through which two computers can communicate in private. A subset of these technologies provides not just privacy but also *redundancy*, and *throughput* guarantees. In addition, they can connect a multitude of locations, such as a bank's branch offices to headquarters, to form what's referred to as a *multi-point network*.

» **Internet protocol security virtual private network (IPsec VPN):** This VPN, implemented using the Internet Protocol Security protocol, authenticates the two endpoints and encrypts the traffic between them. Although the two computers communicate over the Internet and traffic can technically be intercepted, it will be indecipherable. IPsec VPN is typically implemented as a standalone appliance at each location needing VPN. Creating multi-point networks using IPsec VPN is a challenge without a great management tool.

» **Secure Sockets Layer virtual private network (SSL VPN):**

This VPN, implemented using the Secure Sockets Layer (SSL) protocol and in recent implementations SSL's successor, Transport Layer Security (TLS), authenticates the two endpoints and encrypts the traffic between them. Like IPsec VPN, SSL VPN is used over the Internet, and traffic can technically be intercepted, but it will be indecipherable. Unlike IPsec VPN, you can use SSL VPN via standard web browsers, so you don't need dedicated appliances or special software. As such it is referred to as "client-less." Whenever you access a URL beginning with "HTTPS," TLS is being used to encrypt the traffic with no effort or setup by the user.

## Recognizing Factors That Impact Network Performance

Network performance is a constant concern for both network administrators and users. Nobody wants to wait around for web pages or video to load, listen to choppy audio, or have their connections interrupted on a regular basis.

Performance is the product of the following factors:

» **Bandwidth** is the advertised theoretical maximum data transfer rate of a network connection measured in anything from kilobits per second (Kbps) to hundreds of gigabits per second (Gbps). Keep in mind that bandwidth between devices that are linked over two more connections is limited by the slowest connection. If your computer has a 1 Gbps connection to the router, and the router has only a 100 Mbps connection to the Internet, your computer's connection to the Internet is the slower of the two. In addition, many ISPs throttle their bandwidth if a user has exceeded the monthly data allotment.

» **Throughput** is the actual amount of data per unit time that's transferred over a network connection, and it's primarily impacted by latency, discussed next. Generally, the longer the distance data has to travel the lower the throughput will be.

- » **Latency** is the time it takes for a packet to travel from origin to destination on a network; it's typically measured in milliseconds (ms). Latency is primarily due to the physical media being used, the number of different systems each packet has to go through, and TCP/IP, which verifies that data packets have reached their destination. The more time it takes for the sending and receiving devices to confirm receipt, the greater the latency and the lower the throughput.
- » **Resiliency** enables continuity of network service by providing redundant equipment and connectivity throughout the network, coupled with the ability to automatically move network traffic from a non-functional to functional path through *routing*.
- » **Packet loss** occurs when packets travelling across a network are lost or corrupted due to network congestion, network device issues, or connectivity issues.
- » **Jitter** occurs when packets arrive in a different order than they were sent, which can cause mayhem with the applications receiving the data, especially video and voice.
- » **Quality of service (QoS)** refers to the overall performance of a network, relative to specific metrics such as throughput and latency, but most importantly, from the perspective of the users of the network. QoS also can be shorthand for "QoS management," in which network traffic and resources are actively monitored and managed to maintain a certain level of service. QoS controls may be very granular, where specific users, user groups, or applications are given relative priority within the network. Not all networks support the idea of QoS; for example, the Internet treats all users and applications equally.



- » Tracing the history of information technology (IT)
- » Getting your head in the cloud
- » Ushering in a new era of networking with hybrid IT

## Chapter 2

# Exploring Hybrid IT: Setting the Stage for SD-WAN

**M**any businesses are stuck in the past. They've networked their computers and connected them to the Internet, but they're still not reaping the full benefits of technology that's on the cutting edge, such as cloud computing. As a result, they're struggling with issues related to accessibility, performance, and security.

Sometimes the best way to get unstuck from the past is to look back in history to see what brought you to this point. You can then recognize everything you've been holding on to (and perhaps have become overly comfortable with) and then let them go so you can more clearly see what you've been missing out on. In this chapter, we lead you through such an exercise and set the context for the new era of *hybrid IT* — combining on-premises legacy hardware and software with cloud-based applications and resources.



REMEMBER

SD-WAN is a unique approach to hybrid IT. While hybrid IT routes network traffic over two or more connectivity paths, such as MPLS and broadband, SD-WAN uses software to create and manage a virtual network on top of the various physical networks. With this

virtual network overlay, traffic flows through encrypted tunnels, and you control traffic and security across the WAN by setting policies through a single, central interface.

## Understanding How We Got Here: A Brief History of IT

Early networking ushered in an era of increased productivity and reduced costs. By networking computers and other devices, businesses enabled their employees to communicate and share data more efficiently; share expensive equipment such as servers, printers, and high-capacity storage devices; share applications and Internet access; and reduce paperwork.

However, early networks remained relatively isolated. Connecting to the network from a remote location to access files and communicate with coworkers was a major hassle requiring the use of clunky remote-computing utilities, and sharing software and hardware from such a remote connection was impossible.

Information technology (IT) has evolved to address these issues, and it continues to evolve to keep pace with changes in hardware, software, and business models. “State of the art” is a fleeting phenomenon, which keeps IT professionals busy forming strategies and implementation plans to bring the best to their organizations and enhance productivity.

In this section, we describe the characteristics of IT over the past 25 years or so and explain why it’s been that way. This retrospective provides context for the later section on cloud computing.

### Looking back at the early days of centralized computing

Widespread adoption of computing in the '70s and '80s was driven by a computing model that attempted to limit end-user administration and concentrate management centrally. These were the heydays of the mainframe computer, during which users would connect to the mainframe via *dumb terminals* (essentially a keyboard and monitor). Networking speeds were slow and unreliable

in wide area networks (WANs). The result was computing that used end-user applications installed at that site.

In the '90s and early 2000s, client-server computing took off, taking advantage of the power of PCs and reducing the size and cost of the mainframes being used. Mainframes were replaced with smaller servers that hosted individual software applications, further reducing cost.

During this time, WAN speeds picked up and multiprotocol label switching (MPLS) was commercialized, providing a more reliable and secure way to connect locations of an organization. These developments made it practical to consolidate application deployment at organizational headquarters (HQ) in the quest to reduce the IT staff at remote and branch locations.

The result of this consolidation was a network topology that saw branch and remote offices connecting to a hub and spoke model with responsibilities for managing traffic to the outside world (such as the Internet) being concentrated at HQ.

## **Decentralizing computer systems**

As the cost of personal computers (PCs) fell, computer systems became less centralized over the '90s and early 2000s. Every employee had a computer with the software they needed to do their jobs — software developed internally or provided for the mass market, such as Microsoft Word and Excel.

During this era, organizations deployed, maintained, and supported the hardware and software on each user's PC along with any network used to connect the PCs. This required significant technical expertise and often additional personnel. Early on, businesses were concerned with replacing pen and paper with their electronic equivalent, primarily word processors, spreadsheet applications, and databases. Focused on this need, business leaders gave little consideration to connecting with customers and suppliers electronically. However, the advent and widespread adoption of the Internet would soon convince them to change course. To remain competitive, they would need to start conducting at least some of their business on the Internet.

# Opening a New Chapter with Cloud Computing

By now, everyone is familiar with the concept of *cloud computing* — the practice of using a network of remote servers hosted on the Internet, a local server, or PC to use applications and to store, manage, and process data. If you’ve ever used Dropbox to share files or Google Docs to collaborate on documents, you’ve engaged in cloud computing.

However, only a few years ago, few people had ever heard the term. The concept began to pick up steam around 2006, when Google and Amazon started using the term “cloud computing” to describe a growing trend toward accessing storage, processing, and software applications via the Web. In 2006, Amazon launched Amazon Web Services, paving the way for anyone to build and deploy applications in the cloud.

In many ways, cloud computing is an expanded version of using applications on a mainframe computer from a remote terminal. In the cloud, the software application is stored on remote servers, not on individual PCs or a local server. Instead of buying an application and installing it on your PC, you subscribe to the application and access it through your web browser (the application runs on the remote server). This arrangement is referred to as *Software as a Service* (SaaS), in contrast to software as a product. One of the first companies to capitalize on cloud-based computing and SaaS was Salesforce, which pioneered customer relationship management (CRM) SaaS in 1999.



TIP

SaaS delivers numerous benefits:

- » **Low barriers to entry:** Businesses small or large can access the hardware, software, and services they need, paying only for what they use instead of having to buy expensive hardware and software.
- » **Virtually unlimited scalability:** Businesses can easily scale up or down based on their needs, not on limitations inherent in their computer systems.
- » **Multi-tenancy:** Resources are shared, reducing costs.

- » **Device and location flexibility:** Users can access the hardware and software from anywhere using any web-enabled device.
- » **IT management and security outsourcing:** In addition to maintaining and upgrading their hardware and software, SaaS vendors typically employ redundancy to provide consistent availability and performance, and they automatically back up their systems and secure and back up their clients' data.

## Outsourcing IT roles and responsibilities

SaaS and other “aaS” offerings represent a shift in the deployment, operational, and business models for software. Historically, a customer's IT department purchased their software from a vendor (or developed it internally), installed, and maintained the software. The customer purchased annual licenses per seat and had to keep track of these licenses, representing a large annual expense and logistical hassle. Users had no access to the software unless they were on premises.

## NOT JUST SOFTWARE AS A SERVICE

Although SaaS is perhaps the most common cloud computing acronym, you're likely to encounter a host of other “as a service” acronyms, including the following:

- **Platform as a Service (PaaS)** provides the customer with hardware and software usually for application development and deployment. For example, if you want to build a cloud-based application, you would contact a vendor that provides PaaS.
- **Infrastructure as a Service (IaaS)** provides the customer with hardware, network, and physical resources that the customer can access on demand, paying only for the resources used. Unlike SaaS, in which you use applications hosted on the cloud, IaaS enables you to develop and deploy applications that you and others can access from the cloud.

You may also encounter XaaS, which stands for anything and everything as a service.

SaaS fundamentally altered this convention, eliminating most, if not all, of the IT effort and cost related to managing software. The burden of supporting the spectrum of end-user devices is shifted to the SaaS provider, as is the burden of upgrading software and the delay between availability of the latest version and when the end user has access to it. Gone, too, are the servers and storage required to host applications and their data. Expenses shift from capital expenditures to operating expenses and from a large annual fee to pay as you go.

IaaS also represents a significant change in IT infrastructure consumption and the shift away from onsite resources, large capital purchases, and extended depreciation cycles. With IaaS, the costs to deploy, maintain, and administer the hardware and network upon which applications are built and deployed are managed by a third party. Where traditional managed hosting dealt in physical networks and servers, IaaS provides virtual, self-service machines and networks available on demand with pay-as-you-go pricing.

In addition to the infrastructure itself, IaaS providers host a wide range of software that can be incorporated into otherwise custom applications, including database and analytics applications.

## Transitioning from legacy to cloud-native applications

While you can build a new business exclusively using cloud-native applications, existing businesses often have a great deal invested in legacy applications. Here's the difference between the two:

- » **Cloud-native applications:** Software developers no longer need to concern themselves with accommodating the differences that traditionally existed across manufacturers' OS versions and hardware platforms.
- » **Legacy applications:** These applications are typically developed for custom environments that may be specific to server brand and model, operating system levels, code versions, and other specifications that make them inherently non-portable.

Often, the only choice available to the owner of those legacy applications is to live with them as is or completely rewrite them as cloud-native. Transitioning from legacy to cloud-native

applications often takes considerable time, forcing would-be cloud users to live with one foot in the legacy camp and the other in the cloud camp — an approach referred to as hybrid IT, discussed next.

## Hailing the Arrival of Hybrid IT: The Best of Both Worlds

*Hybrid IT* is a combination of traditional on-premises computing with public cloud computing for maximum cost-effectiveness, portability, investment protection, and security. With hybrid IT, organizations don't have to scrap everything they've invested in so far to take advantage of cloud computing. They can keep what they have that's working well for them, and gradually adopt new cloud-based alternatives.

The following examples reveal some of the many ways that businesses are using the cloud to augment their existing on-premises computer systems:

- » **Commercial applications:** Suppose you own a business, most of your employees are already using an on-premises version of Microsoft Exchange, and you want to transition to Microsoft's cloud-based Exchange Online in Office 365. Your on-premises users can continue to use the existing Exchange Server infrastructure while new users and those who want to switch can choose Exchange Online. You can transition with little or no disruption and with all users sharing the same email address space. In addition, you can route all incoming Internet email through Exchange Online Protection, regardless of whether the recipient is using the on-premises or online version of Exchange.
- » **E-commerce:** Exposing software applications to the outside world creates challenges that most businesses are ill-equipped to handle, in terms of both technology and compliance. Further, the scaling that may be necessary to handle workloads is not economical, versus the economies of scale in the public cloud.
- » **Experimental development:** The advent of SaaS not only provides new consumption models for familiar software but

also brings a host of new ready-to-integrate software, such as analytics packages for big-data applications. The expense and expertise required can be prohibitive for private organizations, but in a public cloud setting can become easy to incorporate. This frees up organizations to focus on innovation.


» **Backup and disaster recovery:** Businesses must have a plan to deal with natural disasters, technical calamities, and other unforeseen events. Traditionally, organizations would maintain a secondary location “just in case,” representing a huge expense. Public cloud enables an organization to have a virtual secondary location at their fingertips, without incurring any significant expenses unless otherwise needed. In a related area, the cost of maintaining backups is a huge expense to organizations, which can be mitigated by leveraging the economies of scale enjoyed by cloud providers.

## Recognizing the Challenges Inherent in Hybrid IT

With hybrid IT, an organization’s hardware, software, and data may be distributed among several “nodes,” as shown in Figure 2-1. The organization may have its own network of computers (linked across stores, branches, or offices) that are connected to Microsoft servers for running Office 365 applications, Amazon or Azure servers for IaaS, and one or more data centers. With such an arrangement, you face the challenge of maintaining security, availability, and performance when you have little or no control over the cloud services provided by outside vendors or the connections to those services.

A key challenge is security. The Internet exposes any equipment connected to it to outside threats. Connecting multiple locations together securely has traditionally required expensive private networking connections that took considerable time to get built by telecom companies. The middle ground has been IPsec VPN, which is typically painful to set up and operate and doesn’t solve the availability issue.





© John Wiley & Sons, Inc.

**FIGURE 2-1:** Hybrid IT schematic.

The rapidly growing world of hybrid IT has brought the issue of networking back to the forefront, breaking traditional hub and spoke architectures, rendering traditional service provider options obsolete, and demanding a solution that's easier to manage and align with business policies. The solution to this problem is *software-defined wide area networking* (SD-WAN), as we explain next in Chapter 3.



- » Understanding what SD-WAN is
- » Recognizing SD-WAN's benefits and limitations
- » Comparing SD-WAN to competing technologies
- » Finding a reliable SD-WAN provider

## Chapter 3

# SD-WAN: Policy-Driven Networking

**A**lthough the Internet is a worldwide network of networks, most organizations have to cobble together a collection of hardware, software, vendors, and so on to get all of their internal and external computing resources to function as one unified, reliable, fast, and secure network. Even if they manage to achieve this great feat, they have a host of issues to deal with related to availability, performance, and security. To compound the problem, they have little to no control over external factors, such as Internet service providers (ISPs) and connections between public networks.

To meet these challenges, many organizations are deploying software-defined wide area networking (SD-WAN). In this chapter, we explain what SD-WAN is and how it works, present its benefits and limitations, compare it to traditional approaches, and bring you up to speed on where to look for SD-WAN service providers.

# Introducing SD-WAN

SD-WAN is a technology for managing wide area networks. It's easy to deploy, centralizes network management, improves connectivity, ensures privacy, and reduces costs. The software is either deployed as a stand-alone appliance or as a feature of a next generation firewall (NGFW) at connected locations — company headquarters, branch offices, stores, the public cloud, and so on.

SD-WAN doesn't replace traditional networks but instead creates a new, independent layer on top of them, enabling the collection of networks to operate as a single, unified network. The underlying networks remain the responsibility of the individual network managers. You or an SD-WAN service provider manages the new virtual network through a central orchestrator that provides visibility and control across the entire network. No on-premises expertise is required to get the service up and running.



TIP

A key benefit of SD-WAN is that it enables you to scale bandwidth by aggregating connections from multiple service providers with different connection types (for example, broadband Internet, 4G-LTE, and MPLS), thus improving performance and resiliency while reducing costs. You can select higher quality, higher cost connections for business-critical traffic, and lower quality, lower cost connections for ordinary traffic. In addition, the ability to move data traffic over multiple connections of different types helps to maintain business continuity even if one of the provider's networks goes down.

SD-WAN can keep data private by encrypting the traffic over public WAN connections, which ensures that even if a third party intercepts network traffic, it will be unintelligible to them. This removes the issue of data privacy as a significant barrier to using commercial broadband for one or more of the connections for sensitive data.

The policy management and service monitoring engines of SD-WAN enable it to make decisions on which network links to use to optimize performance for the best end-user experience. It does this by monitoring end-to-end performance, including latency and jitter, on all network links and making decisions on which

networks to use based on the policies defined by the organization. The ability to perform centralized management ensures that policies are defined once and pushed to all network endpoints, ensuring consistency across the entire network.

SD-WAN provides more privacy, control, visibility, and manageability of the wide area connections between the LANs it connects. The result is more predictable performance and reliability, delivering a better user experience and better servicing of the organization's needs. But connecting sites directly to the Internet that previously sent their traffic through a central HQ creates a need for additional network security to prevent attackers from getting into the now-extended corporate network through the new SD-WAN links.

## Recognizing What SD-WAN Does (and Doesn't) Do

SD-WAN improves network speed and accessibility while reducing costs, but it doesn't promise to solve all your WAN-related problems. To help you manage your expectations, in this section, we explore what SD-WAN does and doesn't do.

### It secures the data, but not the network

SD-WAN typically uses encryption technology to shield the data flowing over the WAN from prying eyes, ensuring that sensitive data remains private. In a way, encryption creates tunnels between the LANs that populate the WAN. This removes the requirement for underlying network providers to offer this capability.



REMEMBER

SD-WAN does not replace the need for next-generation firewalls (NGFWs) and intrusion protection services (IPSS), which are still necessary to keep attackers out of LANs. New “Secure Enterprise SD-WAN” products are integrating SD-WAN networking with NGFW security to provide a single, centrally managed way to connect and protect locations together.

## It's resilient

SD-WAN can automatically select from two or more network connections from multiple providers based on link quality and availability, eliminating a single point of failure and ensuring business continuity.

SD-WAN does not manage the underlying networks, nor can it control their performance, because they're managed by third parties and operate using automated protocols.

## It's policy-driven

SD-WAN uses business-driven, user-defined policies to control its behavior and choices. This is a break from the past, where users relied on generic protocol to make decisions on traffic routing. SD-WAN uses multiple connections and intelligently chooses the most suitable path based on user-defined policies. Layering policy upon protocols enables organizations to tailor the behavior of SD-WAN to their specific requirements.

## It reduces costs

Although SD-WAN doesn't eliminate the need for third-party network services, it can save you money in three key areas:

- » **Provisioning:** To ensure sufficient performance at all times, many organizations *overprovision* — they rent bandwidth to meet their peak quality of service (QoS) needs, but that bandwidth goes to waste during off-peak hours.
- » **Transitioning from MPLS:** One of most common reasons organizations adopt SD-WAN solutions is to transition away from older, expensive MPLS lines to less-expensive broadband connections. With SD-WAN, you can send high-priority data over faster connections and low-priority data over slower, less expensive commercial connections.
- » **IT personnel:** With SD-WAN, you centralize WAN management, so you don't have to hire, train, and pay IT professionals at every location to manage the WAN.

## It enhances performance

SD-WAN uses policies, redundant providers, and network performance awareness to keep traffic moving over the most effective path available. SD-WAN can also aggregate Internet connections of different types to create a single bundled connection that's faster than any of the individual connections at a lower cost. You can buy cheaper, slower, commercial-grade Internet connections and bundle them together to boost performance.

However, SD-WAN does not manage or modify the performance of the underlying service provider networks.

## Comparing SD-WAN to Traditional Alternatives

Table 3-1 provides a brief comparison of SD-WAN to traditional WAN services.

**TABLE 3-1** SD-WAN versus Traditional WAN Services

	SD-WAN	IPsec VPN	MPLS	Internet
<b>Private</b>	<u>Yes</u>	<u>Yes</u>	<u>Yes</u> <sup>1</sup>	No <sup>2</sup>
<b>Encrypted</b>	<u>Yes</u>	<u>Yes</u>	No	No
<b>Multipoint</b>	<u>Yes</u>	No	<u>Yes</u>	<u>Yes</u>
<b>Resilient</b>	<u>Yes</u>	No	<u>Yes</u>	No <sup>3</sup>
<b>Policy-driven</b>	<u>Yes</u>	No	No	No
<b>Performance-based routing</b>	<u>Yes</u>	No	No	No

<sup>1</sup>MPLS lines typically were deployed behind organizations' existing network defenses, preventing remote sites from being exposed directly to potential attackers on the Internet.

<sup>2</sup>HTTPS provides security for web pages accessed via web browsers, but it doesn't provide any defense against attackers.

<sup>3</sup>The Internet is resilient, but a single connection to a user's location is not.

## Obtaining SD-WAN

SD-WAN may be obtained as a stand-alone service from traditional telephone companies (telcos) or from new entities whose sole business is to provide SD-WAN service. In both cases, it remains the responsibility of the SD-WAN user to obtain WAN service such as DSL, cable, wireless, or MPLS. Some telcos may resell secondary network services as part of a complete package.

You may wonder why an organization would utilize MPLS as one of the networks for SD-WAN if the goal is to reduce costs. The simple answer is that many companies that have been using MPLS are locked into multi-year contracts or want the performance guarantees offered by MPLS for specific applications. Instead of buying yet more expensive MPLS bandwidth, some organizations will employ SD-WAN to put their most important traffic over MPLS and lesser traffic over broadband.

## The New Trend: Secure Enterprise SD-WAN

SD-WAN is also available as a feature of select next generation firewalls (NGFWs), eliminating the need to purchase SD-WAN as a standalone service and managing SD-WAN as a separate entity. Why firewalls, you ask? Going back to our discussion of hybrid IT (see Chapter 2), organizations need to connect and protect their data and applications regardless of where they reside. SD-WAN provides the connectivity, but by integrating with an NGFW, you eliminate gaps that can arise from having them separate.



- » Appreciating the power of multilink management
- » Distributing traffic for bandwidth efficiency
- » Maintaining surveillance of your SD-WAN network

## Chapter 4

# Putting It All Together: Secure Enterprise SD-WAN

The promise of SD-WAN is to connect users in distant locations to one another and to both organizational and cloud resources via fast, reliable, and secure links. Forcepoint Next Generation Firewall (NGFW) delivers on that promise, while supporting crucial features including clustering, load balancing, and Quality of Service (QoS) that make operations more effective and efficient. (See Chapter 1 for details about clustering, load balancing, and QoS.) In addition, it marries connectivity to security, creating *Secure Enterprise SD-WAN*, which eliminates gaps that can arise when connectivity and security are managed separately. This solution provides a simple and cost-effective way to create secure and reliable high-capacity links between distant locations.

Designed for ease of use, the integrated Forcepoint Security Management Center provides centralized management for configuration at all locations, and it's completely independent of any setup or coordination requirements from the Internet Service Providers (ISPs). The implementation requires no special equipment, licensing, or software and no ISP peering agreements

to ensure high availability. (A *peering agreement* is an arrangement in which ISPs allow traffic from other ISPs to pass through their systems.)

This chapter describes how multilink technology and other features native to Forcepoint NGFW fulfill the promise of Secure Enterprise SD-WAN.

## Introducing Multilink Technology

*Multilink technology* bundles two or more network links to form a single high speed link. Together, the network links deliver capacity and speed. Separately, they ensure availability; if one link goes down, Forcepoint NGFW routes traffic through the other links. All of this happens behind the scenes; the end user notices nothing other than a fast, reliable connection. With multilink technology, organizations can mix and match Internet link types. They can use high-cost, high-speed links for priority traffic and route low-priority traffic over slower, low-cost links. Or, to save even more money, they can bundle multiple lower-cost broadband lines to create a single high-speed link. Faced with the requirement of having always-on connectivity, you can resort to multilink technology to keep your network up and running efficiently and reliably.

In this section we describe various ways to combine several independent ISP links and explain how Forcepoint NGFW enables you to prioritize traffic along the different links. We also provide a couple case studies to demonstrate the power and flexibility of Forcepoint NGFW in real-world implementations.

### Aggregating links

A key feature of SD-WAN is that it enables you to take advantage of *clustering* — bundling computing resources so that they function seamlessly together as a unified resource. In the case of multilink technology, clustering is the ability to aggregate (bundle) Internet links to form a single reliable and high-speed link. Aggregating links delivers benefits in these three areas:

- » **Reliability/availability:** If service is interrupted over one link, traffic is routed over the other links. In addition, random fluctuations in traffic occur at different times over different

links; with multilink technology, users don't experience the fluctuations — they experience only more than acceptable link speeds at all times.

- » **Performance:** Bundling several broadband links creates a single high-speed, high-capacity link. In addition, you can send high-priority data over a high-speed link and lower-priority data over lower-speed links to reduce the impact of traffic congestion on key business operations.
- » **Cost:** Aggregating links reduces costs. Bundling even three commercial broadband links is cheaper than a single leased MPLS link of the same capacity. You also save money by avoiding costly *overprovisioning* — paying for capacity that goes to waste during non-peak hours or paying for additional expensive MPLS capacity to handle non-critical traffic.



TIP

You can bundle Internet links in different ways depending on your needs. While some organizations use two MPLS lines, you can bundle services from different providers using a variety of link types, including fiber, cellular, Metro Ethernet, DSL, and satellite.

## Prioritizing links

Ideally, you want your WAN links to have sufficient networking capacity and no traffic congestion. SD-WAN can achieve this goal through *load-balancing* — distributing traffic over different network links. For example, if you have an MPLS link (high cost) and DSL (cheaper, slower), SD-WAN can direct high-priority traffic over MPLS and lower-priority traffic over DSL. When traffic is light, that lower-priority traffic can travel over MPLS, so you're not wasting that expensive bandwidth.

During peak traffic hours, when congestion occurs, QoS is engaged, either dedicating some of the network links totally to high-priority traffic or guaranteeing a specific percentage of the link capacity to it. Lower-priority traffic has to wait (it can be throttled) or travel over links that aren't reserved for high-priority traffic.

With the preferred link selection provided by Forcepoint NGFW, the QoS functionality allows control over how each application uses the available bandwidth. Mission-critical applications can be placed on links that provide high priority with low latency, while all other applications are placed on the links that have available bandwidth.

Inherent QoS makes the use of network resources more efficient by servicing the most important traffic for your business without your having to purchase more bandwidth.



#### CASE STUDY

Here we examine the case of a global retail company that had been using one MPLS link from each of its locations to its central datacenter, where the main Enterprise Resource Planning (ERP) system was located. Problems arose because the SAP traffic didn't always have enough available bandwidth.

## The problem

The reason for the bandwidth issue was that the other traffic (email, web browsing, and so on) was driven through the same MPLS link. The company wanted to remove that traffic from the MPLS link to ensure that the SAP traffic would always have enough bandwidth.

The company had several offices, so adding a second MPLS link everywhere was too costly. Raising the capacity of the MPLS link was also considered, but it was expensive, and it didn't solve the problem of a single point of failure of the MPLS link.

Even though the company had good service level agreements (SLAs) with its MPLS service provider, the maximum compensation for the link outage was equal only to the subscription fees the company had paid. In the case of a link outage, the compensation wouldn't cover the production losses, so the company wanted to have a cost-effective backup link for the SAP traffic.

## The solution

The retail company solved its problems through the use of Forcepoint NGFW and its Secure Enterprise SD-WAN support. The company purchased a DSL link for all its offices, which is a cost-effective way to supply more bandwidth and backup connectivity to each location.

Forcepoint's multilink technology was used to load balance the traffic between the DSL and MPLS links. The QoS feature was implemented to ensure that SAP traffic always has priority over the MPLS link, and other traffic is automatically directed to the ADSL link. When unused capacity is available on the high-quality MPLS link, the other traffic is able to use it.

In this manner, the expensive and high-quality MPLS link comes close to 100 percent use at all times. At the same time, the cost-effective ADSL link provides capacity expansion and backup connectivity whenever needed.

Here's a sample configuration:

SAP traffic = Priority 1 = Forced over the MPLS link.

HTTP traffic = Priority 4 = Normally over DSL + free capacity on MPLS link.

Using this configuration the company now benefits from guaranteed bandwidth for mission-critical/time-sensitive applications, a better user experience and backup connectivity during any interruptions of MPLS service. Should the DSL link go down, the SAP traffic would have priority over the HTTP traffic.

## Ensuring availability

Organizations often employ redundancy to ensure always-on connectivity. They typically have two MPLS links of equal capacity, so if service is interrupted over one link, traffic can be diverted to the other. Forcepoint NGFW supports redundancy by enabling you to bundle links. Even better, it automatically diverts traffic to other links when one of your links is interrupted for any reason.



CASE STUDY

An industrial organization had their production sites in the United States and one of their sales offices in Bermuda.

### The problem

The Bermuda office depended entirely on the link to the company's production sites. They had only one MPLS link between the production site and the Bermuda office. The CIO felt anxious because Bermuda is a known hurricane area. One big hurricane could disrupt the communication lines and put the company out of business for a long time.

The company compared several options, including satellite backup links and an additional MPLS link from another service provider. All alternatives turned out to be rather complex and costly.

## The solution

The company solved the problem by using Forcepoint's NGFW solution with two MPLS links from different service providers. This configuration enabled them to avoid complex setup and routing configuration between the two MPLS carriers and gain highly available links.

About one year after implementing Forcepoint NGFW, a category four hurricane swept through Bermuda and took down one of the main service providers. When that event hit the news, Forcepoint support personnel called the organization's IT manager and asked if he'd noticed that one of the company's service provider's networks had gone down. The IT manager said that he hadn't noticed anything — traffic was flowing flawlessly. This is just one example of how Forcepoint's Secure Enterprise SD-WAN capability overcomes the challenge of availability.

## Load Balancing Multilink Connections

Load balancing is often described as “distributing traffic evenly over network links,” but it involves more than that. Ideally, you want as much traffic as possible traveling over your highest speed link without congestion. When traffic is light, you want it all flowing over your highest-speed link. When it's heavy, you want important stuff flowing over the highest speed link and less important stuff taking other routes. If your load balancer isn't properly configured, you're likely to encounter problems, such as the following:

- » Traffic goes to only one service provider link, even though multiple active links are available.
- » Traffic goes to a poor-quality link, even though a better link is available.
- » Traffic goes to a standby link, even though an active link works.
- » Switching to a standby link takes too long.



REMEMBER

Forcepoint NGFW doesn't require a separate load balancer. Load balancing is built in.

# Securing SD-WAN Networks

SD-WAN networks are usually “private” because they encrypt network traffic between locations. The encryption process converts the data being transmitted across the network, which may be human-readable text (usernames, passwords, and the like) and converts it to ciphertext, which is unreadable. Even if SD-WAN traffic is intercepted by a third party, it can’t be read, which ensures privacy. But this alone does not create “secure SD-WAN.”

Standalone SD-WAN solutions don’t and can’t screen the sources of network traffic and so will accept any network traffic that’s routed to it. They also do not inspect the traffic that passes over the links created by SD-WAN. These two jobs are best handled by NGFWs. NGFWs regulate what traffic is allowed to pass in and out of organizations, who is allowed to send and receive that traffic, and what applications may be used to do this. NGFWs also screen traffic for malicious payloads and prevent them from reaching their destination.

Therefore, SD-WAN must work hand in hand with NGFWs to ensure only valid communications take place over SD-WAN infrastructure and protect the security of organizations that use it. It only makes sense to build SD-WAN functionality directly into NGFWs to ensure that security business policies are consistent between the firewall and network functions as well as to provide the most efficient management possible, all through a single pane of glass.

## Taking the Next Steps

With your newfound knowledge, embarking on a Secure Enterprise SD-WAN deployment is straightforward. As you prepare to deploy SD-WAN, complete the following checklist:

- » Identify the business locations you need to connect, including third-party clouds.
- » Identify all applications and data in remote locations that must be secured and determine the relative importance of each application and data source.

- »» Decide whether your organization needs a stand-alone SD-WAN service or SD-WAN built into the NGFW.
- »» Decide who will manage the organization's SD-WAN.
- »» Estimate the total bandwidth required for each location, keeping in mind that you can aggregate bandwidth from two or more providers.
- »» Select two (or more) broadband providers to service each of the locations your organization wants to connect. (For MPLS customers, one provider may be sufficient.) Consider how flexible each provider is if bandwidth needs to scale over time.



**REMEMBER**

SD-WAN is used in conjunction with, not instead of, broadband services.



- » Recognizing what your current WAN is lacking
- » Matching your needs to solutions
- » Highlighting the need for strategic planning

# Chapter 5

## Ten Questions to Evaluate Your Needs

In this chapter we present a list of ten questions to ask yourself about your company. Answering this list of questions and referring to the other relevant parts of this book help you start to tap the benefits of Secure Enterprise SD-WAN and give you a framework to help you quickly grasp the situation you face with your network today and then build your strategies accordingly.

### What Role Does the Network Play?

To get a clear idea of your networking needs, examine your business closely in the following three ways:

- » Define exactly what the network brings to you in terms of business value and which activities depend on it most.
- » Consider the consequences of downtime, broken connections, stealth attacks, or increasing traffic volumes.
- » Consider your cloud computing needs or the opportunities the cloud provides that you may not be taking advantage of.

## How is My Network Connected?

Here you need to know exactly what you subscribe to and why: digital subscriber lines (DSL), leased lines, cable modems, satellite, mobile broadband, and perhaps even WAN links such as point-to-point MPLS. Knowing the full scope of links you use and the amount of bandwidth you're currently consuming influences the decisions you make regarding service providers, bandwidth requirements, and network configuration.

## Do I Need a Service Level Agreement?

Obtaining a clear picture of the quality of the connections you use helps you decide whether you need to change ISPs or offerings, implement link-balancing technology, do both, or, in a best-case scenario, do nothing. How great would that be! (See Chapter 1 for a brief description of SLAs.)

## How Available Are My Applications?

You already know which applications are critical for your business. Take that knowledge one step further and analyze whether your current networking solutions are fully supporting application availability or hindering it. Use NGFW's cloud application discovery tool to see how much your organization is already using cloud applications. If your current solutions are restricting application availability, you can start to find changes to help here.

## Is the Transmitted Data Secure?

Cyber security is one of the highest concerns for governments and business today, and the fact that security precautions aren't always implemented is becoming increasingly clear. Taking care of your security helps you to take care of your business and avoid the disaster of breaches and non-compliance.

See Chapter 1 to find out more about encryption and VPN protocols that you need to use to secure your network connectivity.

## Am I Equipped for Increased Traffic?

Even if your traffic flows are incident-free at any given time, you need enough foresight to prevent bottlenecks before they start becoming a persistent issue. Accessing an application from newly opened sites or lacking backup solutions in the event of an outage can quickly lead to significant losses in productivity.

Find out how Forcepoint's multilink technology enables you to scale bandwidth to optimize network performance in Chapter 4.

## Are Setup and Updates Taking Too Long?

Using traditional WAN solutions, organizations typically spend an excessive amount of time on preliminary firewall configuration, policy setting and updates, remote location setup, and more. Research commissioned by Forcepoint reveals a significant divide between the ideal and the real investment in IT management.

Although you may not be able to predict today how much time and effort you'll spend on future installations and configurations, you can easily imagine the time and effort you'll save through centralized remote site management.

Chapter 4 reveals a solution that's been proven to slash time spent on WAN setup and configuration an average of 70 percent.

## Am I Using the Fastest Connection?

You may sometimes receive reports that an application took longer than usual to access, or that recipients didn't see the information intended for them. Forcepoint NGFW helps to ensure that you're always using the fastest connection in two ways:

- » Multilink technology automatically directs network traffic to give high-priority traffic access to the fastest connection. Chapter 4 presents the different methods that Multi-link technology uses to select the fastest available connection.

- » Security Management Center provides complete single-pane-of-glass visibility of physical and virtual networks. You can examine a history across all the links you deploy and use the data to identify the source of bottlenecks.

## Do I Put Performance Over Security?

As security becomes more and more complex, a tug-of-war has emerged, with network administrators facing situations where advanced protection can adversely affect network performance. Unfortunately, they often choose performance over security, a choice that organizations should not be forced to make. For example, turning off deep packet inspection to accelerate traffic flow will work well to decrease latency — until the day an attack slips by unnoticed.

You should therefore ensure that you obtain the right balance between full security features and optimum throughput. Examine security and performance issues across your network, so you're not slowing traffic with unnecessary security and not exposing traffic to security breaches in an effort to boost performance.

## Can I See All Connections in Real Time?

Whether you speak of electrical circuits, railway infrastructures, or digital connections, a global overview of what's happening on any network at each instant is undoubtedly precious. With distributed networks, the ability to supervise from a central point and from anywhere in the world is indispensable.

Browse through Chapter 4 for insight into SD-WAN management and surveillance.

- » Cutting costs in more ways than one
- » Keeping pace with the speed of business
- » Streamlining WAN deployment and management

## Chapter 6

# Ten Advantages of Secure Enterprise SD-WAN

One of the biggest myths about SD-WAN is that it's *all* about replacing private MPLS networking links with direct-to-Internet connectivity over commodity broadband to save money. While SD-WAN can certainly replace MPLS with commodity broadband, and doing so does save money, that's not the *only thing* it's about.

Secure Enterprise SD-WAN cuts costs in more ways than one and delivers far more benefits than just providing an alternative to the tangle of MPLS lines. In this chapter, we reveal ten advantages of Secure Enterprise SD-WAN, so you have a better appreciation of all it has to offer.

## Cutting Costs

Yes, SD-WAN cuts costs, but it does so in several ways:

- » **Leverages the cost savings of commodity broadband.**  
With SD-WAN, you can aggregate connections from

numerous broadband providers to improve network connection speeds and reliability while cutting costs significantly.

» **Reduces hardware investment.** Improved connectivity enables organizations to make greater use of cloud applications, reducing the need for on-premises hardware. You spend less on hardware and on maintenance.

» **Reduces IT costs and hassles:**

- SD-WAN is managed centrally, so you don't need to hire IT personnel at every location to manage the WAN.
- Eliminates *overprovisioning* — renting bandwidth to meet your peak quality of service (QoS) needs, and having it go to waste during off-peak hours.
- The increased use of cloud applications eliminates the hassles of managing software licenses, the need to install updates and upgrades and the troubleshooting to resolve compatibility issues.

## Replacing MPLS with Commodity Broadband

The most touted benefit of SD-WAN is that it enables an organization to replace MPLS connections, often perceived to be expensive and slow, with lower cost commodity broadband. SD-WAN enables the aggregation of broadband connections, so several slower connections function as a single high-capacity connection.

One customer described it as a “10x choice” replacing MPLS lines with commodity broadband at one tenth the cost. However, they saw as much, or even more, value in going the other direction — increasing capacity tenfold with the same budget.



REMEMBER

In many cases, MPLS lines aren't immediately removed, either because they're part of multi-year subscription contracts or are carrying sensitive traffic that would require additional business processes if sent over external links. SD-WAN gives organizations the flexibility to augment their MPLS lines with other types of connections as new sites are deployed or as existing sites require greater capacity.

## Enhancing Business Continuity

With SD-WAN, multiple links can be used together seamlessly so that when one connection goes down, the network may slow down, but it doesn't shut down. Organizations that rely on a single high-capacity MPLS connection are at a greater risk of business disruptions due to accidents and catastrophic events, such as a networking cable that gets dug up outside the building or a hurricane that takes out an MPLS supplier.

## Accessing High-Productivity SaaS Applications

Newer, highly interactive cloud applications such as Office 365 depend on users being connected as directly as possible to the Internet. SD-WAN enables sites to connect to the Internet without going through intermediate systems back in central offices. Organizations can use cloud applications more confidently, without the threat of slow or broken network connections interrupting work.

In other words, one of the benefits of SD-WAN is that it enables you to more easily reap the benefits of cloud applications, such as:

- » Automatic software updates
- » Built-in disaster recovery
- » Eliminations of capital expenditures
- » Increased mobility (you can access the applications from any Internet-enabled device)
- » Enhanced scalability
- » Enhanced collaboration

## Isolating Sensitive Data to Reduce Risks

As compliance mandates (and the audits that check up on them) become more complex, organizations often find that segmenting their networks so that sensitive data (such as PCI-controlled

financial information) is kept separate from other more “general” network traffic can substantially reduce their risk (and make their auditors happier).

## Optimizing Infrastructure

MPLS and certain other types of links provide reliable quality of service for applications such as voice-over-IP (VoIP) that need predictable delivery. But using such links for all traffic can be overkill and expensive. With SD-WAN, each application can use the most appropriate network connection, providing better overall performance and cost at the same time.

## Keeping Attackers Out

Industry-leading analyst Gartner Group advises that branch locations require the same level of security that organizations employ at their primary Internet gateway. Having full NGFW security that’s tightly integrated with SD-WAN connectivity prevents gaps from forming that attackers could exploit to slip into the enterprise.

## Being Enterprise-Ready

Not all SD-WAN solutions are enterprise-ready, providing the high availability, manageability, and security that highly distributed organizations require. True enterprise-grade SD-WAN is designed to:

- » Connect more than a thousand locations over whatever network links are appropriate for each site.
- » Provide always-on resilience for 24x7x365 productivity.
- » Tightly integrate security with connectivity to prevent gaps.
- » Provide full visibility of user behaviors and network traffic across the WAN.
- » Enable common policies to be expressed once and reused automatically wherever needed.



- » Enable special requirements at different locations to be expressed efficiently without disrupting normal operations.
- » Automate connectivity tasks such as setting up VPNs among sites so that new sites can be added quickly and reliably.
- » Dynamically push policies to modify how sites are connected or secured.
- » Update the networking and security infrastructure without taking sites offline.



WARNING

Don't settle for less than truly secure enterprise-grade SD-WAN. Many SD-WAN vendors can only point to deployments with dozens or perhaps a hundred locations. As a result, they often have management systems that are geared towards single sites or relatively small numbers, or they assume that all sites have the same types of connections.

## Boosting WAN Performance

SD-WAN keeps traffic moving over the most effective paths available. You can buy cheaper, commercial-grade Internet connections and bundle them together to boost capacity while lowering costs.

## Increasing Visibility across the WAN

Forcepoint NGFW's Security Management Center provides 360-degree visibility into users' behaviors and the flow of data everywhere, from branches to main offices and from data centers to the cloud. It enables you to know what's happening across your entire network, eliminating blind spots that can lead to problems.

## Simplifying Deployment

As explained in the earlier section "Cutting Costs," SD-WAN provides centralized WAN management, so you don't need IT personnel at each branch to deploy and manage the WAN. Using Forcepoint NGFW's Security Management Center, you can add a new network to the WAN with just a few clicks.



# Use SD-WAN to connect and protect your business

Cost, speed, and reliability are all important to consider when designing a network, but balancing all three is no easy task. As your organization expands, you need a resilient, fast, and secure network — especially in large, distributed environments. And it needs to be affordable.

This is where a software-defined WAN (SD-WAN) comes to the rescue. And deploying it in enterprises today can provide a quantifiable competitive advantage. Gartner estimates SD-WAN has less than 5% market share today, but predicts **25% of users will manage their WAN through software within two years.**

This book provides step-by-step guidance on how to manage and secure digital networks in your enterprise with SD-WAN, with instructions even the most novice networking professional can understand.

## Inside...

- Find networking technology updates
- Gain the terminology you need to know
- Find out how hybrid IT is changing
- Prepare for the cloud era
- Understand SD-WAN pros and cons
- Read real-world SD-WAN case studies
- Learn how to get started with SD-WAN



**Joe Kraynak** ([JoeKraynak.com](http://JoeKraynak.com)) is a veteran Dummies writer who has authored and co-authored dozens of books on a variety of topics.

**Go to [Dummies.com](http://Dummies.com)**<sup>®</sup>  
for videos, step-by-step photos,  
how-to articles, or to shop!

**for  
dummies**<sup>®</sup>  
A Wiley Brand



Also available  
as an e-book

ISBN: 978-1-119-50927-1  
Not for resale

# **WILEY END USER LICENSE AGREEMENT**

Go to [www.wiley.com/go/eula](http://www.wiley.com/go/eula) to access Wiley's ebook EULA.