

LEARNING MADE EASY

Palo Alto Networks Special Edition

Secure Access Service Edge (SASE)

for
dummies[®]
A Wiley Brand



Reduce networking
and security complexity

—
Stop cyberattacks with
consistent security

—
Increase business
speed and agility

Brought to
you by



Lawrence Miller

About Palo Alto Networks

Palo Alto Networks, the global cybersecurity leader, is shaping the cloud-centric future with technology that is transforming the way people and organizations operate. Our mission is to be the cybersecurity partner of choice, protecting our digital way of life. We help address the world's greatest security challenges with continuous innovation that seizes the latest breakthroughs in artificial intelligence, analytics, automation, and orchestration. By delivering an integrated platform and empowering a growing ecosystem of partners, we are at the forefront of protecting tens of thousands of organizations across clouds, networks, and mobile devices. Our vision is a world where each day is safer and more secure than the one before. For more information, visit www.paloaltonetworks.com.



Secure Access Service Edge (SASE)

Palo Alto Networks Special Edition

by Lawrence Miller

**for
dummies**[®]
A Wiley Brand

Secure Access Service Edge (SASE) For Dummies®, Palo Alto Networks Special Edition

Published by
John Wiley & Sons, Inc.
111 River St.
Hoboken, NJ 07030-5774
www.wiley.com

Copyright © 2020 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, or visit www.wiley.com/go/custompub. For information about licensing the *For Dummies* brand for products or services, contact BrandedRights&Licenses@Wiley.com.

ISBN 978-1-119-69602-5 (pbk); ISBN 978-1-119-69608-7 (ebk)

Manufactured in the United States of America

10 9 8 7 6 5 4 3 2 1

Publisher's Acknowledgments

Some of the people who helped bring this book to market include the following:

Project Editor: Elizabeth Kuball
Associate Publisher: Katie Mohr
Editorial Manager: Rev Mengle
Business Development Representative: Karen Hattan

Production Editor:
Tamilmani Varadharaj
Special Help: Shannon Bonfiglio,
Carmine Clementelli,
Matt De Vincentis,
Rachna Srivastava

Table of Contents

INTRODUCTION	1
About This Book	2
Foolish Assumptions	2
Icons Used in This Book	3
Beyond the Book	3
CHAPTER 1: The Evolution of Networking	5
The Journey to the Cloud — And Beyond	5
The Rise of Mobile Computing	6
The Impact on Branch Networking and WAN Architectures	8
The SASE Vision	10
Modern Networking and Security Challenges Revisited with SASE	11
CHAPTER 2: SASE Use Cases	13
Mobile and Remote Users	13
The limitations of traditional remote access VPNs	14
Unsatisfactory compromises	15
A modern architecture for the mobile workforce	16
Branch and Retail	17
The challenges of traditional branch/retail networking	18
Augmenting MPLS with direct Internet access	19
A modern architecture for branch transformation	20
CHAPTER 3: SASE Networking Capabilities	21
Software-Defined Wide-Area Network	21
Virtual Private Network	25
Zero Trust Network Access	27
Quality of Service	29
CHAPTER 4: SASE Security Capabilities	31
Firewall as a Service	31
Domain Name System Security	33
Threat Prevention	36
Secure Web Gateway	37
Data Loss Prevention	40
Cloud Access Security Broker	42

CHAPTER 5: Ten Benefits of SASE	47
Complete Visibility across Hybrid Environments	47
Control of Users, Data, and Apps	48
Monitoring and Reporting.....	48
Less Complexity.....	49
Consistent Data Protection Everywhere.....	49
Reducing Costs	50
Lower Administrative Time and Effort.....	50
Reducing Need for Integration	50
Better Network Performance and Reliability	51
Greater Agility	51
GLOSSARY	53

Introduction

With increasing numbers of mobile users, branch offices, data, and services located outside the protections of traditional network security appliances, organizations are struggling to keep pace and ensure the security, privacy, and integrity of their networks and, more important, their customers.

Today, many of the current network security technologies on the market were not designed to handle all of the types of traffic and security threats that a modern organization has to deal with. This forces organizations to adopt multiple point products to handle different requirements, such as secure web gateways, firewalls, secure virtual private network (VPN) remote access, and software-defined wide area networks (SD-WANs). For every product, there is an architecture to deploy, a set of policies to configure, an interface to manage, as well as its own set of logs. This creates an administrative burden that introduces cost, complexity, and gaps in security posture.

To address these challenges, secure access service edge (SASE) has emerged. Originally defined by Gartner, a SASE (pronounced “sassy”) solution is designed to help organizations embrace cloud and mobility by providing network and network security services from a common cloud-delivered architecture. A SASE solution must provide consistent security services and access to all types of cloud applications — for example, public cloud, private cloud, and software as a service (SaaS), platform as a service (PaaS), and infrastructure as a service (IaaS) — delivered through a common framework. By removing multiple point products and adopting a single cloud-delivered SASE solution, organizations can reduce complexity while saving significant technical, human, and financial resources.

In *SASE For Dummies*, you’ll learn about this new approach to networking and security, including its core capabilities and key benefits for organizations in the modern digital workplace.

About This Book

Secure Access Service Edge (SASE) For Dummies consists of five chapters that explore the following:

- » Modern trends and their impact on the evolution of networking architectures (Chapter 1)
- » SASE use cases (Chapter 2)
- » Networking capabilities in SASE (Chapter 3)
- » Security capabilities in SASE (Chapter 4)
- » Key SASE benefits (Chapter 5)

Each chapter is written to stand on its own, so if you see a topic that piques your interest feel free to jump ahead to that chapter. You can read this book in any order that suits you (though I don't recommend upside down or backward).

There's also a glossary in case you get stumped on any acronyms or terms.

Foolish Assumptions

It's been said that most assumptions have outlived their usefulness, but I assume a few things nonetheless!

Mainly, I assume that you work in an organization that is looking for a better way to simplify your approach to networking and security services. Perhaps you're an IT executive or manager such as a chief information officer (CIO), chief technology officer (CTO), or chief information security officer (CISO). Or perhaps you're a network or security architect or engineer.

As such, this book is written for technical readers with a general understanding of cloud, networking, and security concepts and technologies.

If any of these assumptions describes you, then this is the book for you. If none of these assumptions describes you, keep reading anyway — it's a great book and you'll learn quite a bit about SASE.

Icons Used in This Book

Throughout this book, I occasionally use special icons to call attention to important information. Here's what to expect:



REMEMBER

This icon points out important information you should commit to your nonvolatile memory, your gray matter, or your noggin — along with anniversaries and birthdays!



TIP

Tips are appreciated, never expected — and I sure hope you'll appreciate these useful nuggets of information.



WARNING

These alerts point out the stuff your mother warned you about (well, probably not), but they do offer practical advice to help you avoid potentially costly or frustrating mistakes.

Beyond the Book

There's only so much I can cover in 64 short pages, so if you find yourself at the end of this book, thinking, "Gosh, this was an amazing book! Where can I learn more?" check out www.paloaltonetworks.com/prisma/access.

IN THIS CHAPTER

- » Understanding the role of the cloud in digital transformation strategies
- » Untethering users with mobile computing
- » Evolving the network architecture
- » Discovering a new approach to enterprise networking and security
- » Addressing networking and security challenges with SASE

Chapter 1

The Evolution of Networking

In this chapter, you learn how cloud and mobile computing trends have changed enterprise networking and how a secure access service edge (SASE, pronounced “sassy”) can help your organization address its modern networking and security requirements.

The Journey to the Cloud — And Beyond

We live in an age of cloud and digital transformation. Users and applications are moving outside the traditional network perimeter, accessing an ever-increasing number of applications — including software as a service (SaaS), platform as a service (PaaS), and infrastructure as a service (IaaS) application workloads in the public cloud. Organizations face the challenge of proactively protecting their users, applications, and data from security threats, without compromising user experience.

The 2019 *RightScale State of the Cloud Report* from Flexera found that public cloud adoption among organizations has grown to 91 percent and companies now run a majority of their workloads in the cloud (38 percent of their workloads run in public cloud and 41 percent run in private cloud). Companies are also using SaaS, PaaS, and IaaS offerings from multiple cloud providers — nearly five clouds on average.

As cloud computing continues to play an integral role in digital transformation, the enterprise network must evolve to support new technologies and business initiatives.

The Rise of Mobile Computing

The proliferation of mobile devices in our everyday lives is indisputable. According to the June 2019 *Ericsson Mobility Report*, there are now nearly 8 billion mobile subscriptions worldwide. By the end of 2024, Ericsson predicts that 95 percent of all subscriptions will be mobile broadband. Many smartphones now contain more computing power than the average desktop computer. People are increasingly using smartphones to access the Internet and SaaS apps, not only for personal computing needs, but also for work purposes.

At the same time, public Wi-Fi hotspots are now available practically everywhere. This ubiquitous connectivity enables users to work on their laptops, tablets, and smartphones from practically anywhere.

Organizations are increasingly taking advantage of these trends by implementing bring your own device (BYOD) policies and embracing remote working as a new norm in the modern digital workplace. Remote working increases productivity and, ironically, promotes a work-life balance that many employees prefer instead of commuting to an office and “clocking in and out” every day.

Mobile and remote computing introduce new networking and security challenges that traditional remote access connectivity is not designed to address.

THE TOP FIVE MOBILE SECURITY THREATS

Mobile devices have emerged in recent years as the leading platform for cybercrime and cybersecurity threats against organizations. However, organizations are still working on ways to protect these mobile devices, especially because they often contain a mixture of business and personal data and operate both on and off the enterprise network.

Here are some of the top mobile security threats today:

- **Phishing:** In the past, phishing attacks largely took place by email. Today, they're primarily happening through mobile channels, such as text messaging, Facebook Messenger, WhatsApp, and phony websites that look legitimate.
- **Mobile malware:** Every website visited or link clicked has the potential to infect mobile devices with malware, such as spyware, ransomware, Trojan viruses, adware, and others. This risk of infection by malware on mobile devices is often higher than on desktop or laptop computers because most mobile users don't install anti-malware software on their smartphones and tablets and don't recognize the threat.
- **Fake public Wi-Fi networks:** Many mobile workers today use public Wi-Fi networks at coffee shops, airports, restaurants, and other locations whenever they're working outside the office. Cybercriminals are aware of this trend and often leverage these networks to trick mobile users into connecting to fake Wi-Fi networks, thereby potentially compromising sensitive data.
- **Malicious apps:** The world is full of software applications that can be used over the Internet or downloaded from websites (including the Apple App Store and Google Play Store). Many of these applications are legitimate and safe to use, but there are also thousands that aren't. Thus, downloading an app or granting an app permission to access functions on a mobile device may expose the user's company to a host of security and privacy risks. Some apps even collect data without asking the user for permission.

(continued)

(continued)

- **Data leaks:** Data leaks occur with any unauthorized or unintentional transfer of data from inside an organization to an external party or destination. These leaks can range from someone inside a company accidentally transferring confidential or sensitive data to an unsanctioned/unapproved cloud application or oversharing confidential or sensitive data on cloud sharing apps or public cloud storage, all the way to an attacker or a disgruntled employee deliberately stealing the company's data. Mobile devices, which often contain both business and personal data, make it even easier to blur the boundaries either inadvertently or maliciously.

The Impact on Branch Networking and WAN Architectures

In the early 2000s, multiprotocol label switching (MPLS) networks began to replace traditional asynchronous transfer mode (ATM) and private leased line hub-and-spoke WAN architectures. Over the next decade, MPLS became the prevalent enterprise WAN architecture. MPLS networks provided a simple network connection between branch offices and central headquarters or data center sites. This design worked well because, at the time, most network traffic was between client desktop computers located in headquarters and branch offices and business applications hosted on servers in the on-premises data center. Internet traffic volume was relatively low and generally consisted of email and static web page browsing. Any Internet-bound traffic — including traffic from the branch offices, which traversed the MPLS connection to the central headquarters or data center sites — was sent through the perimeter firewall for security protection. All network traffic could be inspected, and a centralized security policy could be enforced by the perimeter firewall.

As Internet usage increased, many branch offices began to experience performance issues and latency as their Internet traffic was being backhauled across the MPLS connection and inspected by the perimeter firewall, which was becoming a bottleneck. The growing congestion on the MPLS network negatively impacted both Internet traffic and data center traffic. The rapid adoption of cloud-based SaaS applications amplified this problem

exponentially and essentially put the final nail in the MPLS coffin. Organizations began to provision direct Internet access (DIA) connections, such as broadband, for their branch offices from local Internet service providers (ISPs) to alleviate some of this congestion.

Adding DIA connections at branch offices alleviated some of the network congestion issues but introduced a whole new set of challenges. On the networking side, these challenges include

- » **Routing complexity:** Routers need to be configured to send traffic over the appropriate network link (for example, data center traffic over the MPLS link and Internet traffic over the DIA link). In most cases, the simplest solution is to configure static routes, which provide only limited resiliency.
- » **Inefficient bandwidth usage:** It may be possible in certain cases to configure some basic round-robin load balancing between multiple Internet connections, but more advanced algorithms that take distance, cost, load, or other weighted factors into account are generally not available. As a result, there may be times when the DIA link is congested, but the MPLS link — which could otherwise be used to backhaul Internet traffic through the headquarters or data center Internet connection — is relatively idle.
- » **Management complexity:** In many cases the local Internet service provider (ISP) provides a commodity router for the DIA link and does not give the customer management access. Even if the customer has management access, the ISP routers likely won't be the same type as the MPLS routers. This means different management interfaces, different operating systems, and different remote administration tools — multiplied by the number of different remote locations, different ISPs, and different router models that you need to manage.

On the security side, challenges created by this evolved WAN architecture include

- » **Loss of visibility and control:** With most network traffic traversing the DIA connection at remote offices destined for the cloud and the Internet, enterprise security teams are no longer able to see the traffic and apply security policies from a centralized perimeter firewall in the data center, thereby significantly increasing risk.

- » **Lack of integration and interoperability:** To address the loss of visibility and control, many organizations deploy firewalls, intrusion prevention systems (IPSs), web content filters, data loss prevention (DLP), and other point security solutions in their remote offices. These solutions often come from different vendors and have only limited or no integration capabilities. This makes it more difficult for security teams to correlate events and implement a cohesive enterprise security strategy.
- » **Management complexity:** Different security solutions from different vendors means different management interfaces, different operating systems, and different remote administration tools — multiplied by the number of different remote locations that you need to manage. This management complexity challenge is exponentially more difficult on the security side (compared to the networking side), because of the volume and types of security information that must be analyzed on a daily basis from these different tools.

The SASE Vision

In order to address the shift in networking and security requirements, a new architecture is needed. Gartner writes about a model known as the *secure access service edge* (SASE, pronounced “sassy”).

A SASE solution converges networking and security services into one unified, cloud-delivered solution (see Figure 1-1) that includes the following:

- » **Networking**
 - Software-defined wide area networks (SD-WANs)
 - Virtual private networks (VPNs)
 - Zero Trust network access (ZTNA)
 - Quality of service (QoS)
- » **Security**
 - Firewall as a service (FWaaS)
 - Domain Name System (DNS) security

- Threat prevention
- Secure web gateway (SWG)
- Data loss prevention (DLP)
- Cloud access security broker (CASB)




FIGURE 1-1: SASE delivers advanced network and security capabilities in a converged cloud-delivered solution.

Modern Networking and Security Challenges Revisited with SASE

With networking and security functions unified in a single, multifunction cloud-delivered solution, the challenges of modern networking and security are solved by SASE in the following ways:

- » **Lower capital costs:** SASE requires relatively lower capital investments than other approaches. SASE delivers networking and security capabilities in the cloud, with minimal hardware or software required on-site or on the users' device.

- » **Full visibility and control:** SASE provides full visibility and control with cloud-delivered capabilities including FWaaS, SWG, DLP, and SaaS security via CASB functionality.
- » **Less complexity:** All management functions for the cloud service can be centrally managed in the cloud from an intuitive single-pane-of-glass management interface. This means network and security teams no longer need to learn, configure, and manage multiple systems from different vendors.

- » Enabling mobile and remote users
- » Connecting and securing branch and retail locations

Chapter 2

SASE Use Cases

In this chapter, you learn about some of the most common use cases today for a secure access service edge (SASE).

Mobile and Remote Users

Securing mobile users with traditional types of network security can be a challenge, especially when users work in areas where you don't have IT staff or it's cost-prohibitive to have IT staff in many locations. For years, the standard tool for connecting mobile users into a corporate network was remote-access virtual private networks (VPNs). In fact, for many people, *remote access* and *VPN* are synonymous.

However, with the number of applications and workloads moving to the cloud, the need for remote access is diminishing. In addition, it's apparent that organizations need more than remote access — they need secure access to cloud applications and the Internet as well.

The limitations of traditional remote access VPNs

Remote access VPNs are primarily built to do one thing: Allow users outside the perimeter firewall to access resources inside the corporate network.

Remote-access VPNs use a hub-and-spoke architecture (see Figure 2-1), with users connected by tunnels of various lengths depending on their distance from the data center. Nearby users may enjoy high performance, but distance degrades performance, introducing issues with speed, bandwidth, and latency. Nevertheless, this is the optimal architecture for data center applications because the goal is to reach the “hub” where your internal applications and data are located.




FIGURE 2-1: Traditional remote-access VPN architecture.

The model breaks down when a mixture of cloud applications is involved. With remote-access VPN, traffic always goes to the VPN gateway first, even if the application is hosted in the cloud (see Figure 2-2). As a result, the traffic goes to the VPN gateway at the corporate headquarters or data center and then egresses from the perimeter firewall to the Internet, with the application response going back to headquarters or the data center before it returns to the user. With cloud applications, this traffic essentially follows a “trombone” path, making a lengthy (and slow!) trip to reach an Internet-accessible location. This is sensible from a security perspective, but it doesn’t make sense for network optimization.




FIGURE 2-2: Traditional remote-access VPN backhauling traffic to reach the cloud.

Using cloud applications over remote-access VPN can hurt the user experience, and as a result, end users tend to avoid using remote-access VPN whenever possible. They tend to connect when they need access to the internal data center and disconnect when they don't, which leads to multiple issues. When users are disconnected, their organizations lose visibility into application usage, control over access to unsanctioned applications, and the ability to enforce security policies.

Unsatisfactory compromises

To compensate for the networking problems with remote-access VPN, IT teams typically introduce a number of compromises with certain security implications:

- » **User-initiated tunnel:** A common remote-access VPN deployment model is to let users initiate the tunnel as needed. Typically, they connect for a short time, complete their work with a given application, and disconnect. When disconnected, they have direct access to the Internet with no traffic inspection.
- » **Split-tunnel VPN:** A common yet insecure method of deploying remote access VPN is to set up a policy that permits split tunneling. In this model, traffic bound for the corporate domain goes over the VPN tunnel, and everything else goes directly to the Internet. The improvements in network performance come at a cost, though: There is no traffic inspection for Internet and cloud traffic.

» **Web proxy/secure web gateway:** To compensate for scenarios when users are not connected to the VPN, many organizations have tried alternative network security measures, such as using a proxy for the web browser when users are off-network. However, by definition, a web proxy doesn't fully inspect network traffic. Even worse, the traffic inspection the proxy does perform will be fundamentally different from the inspection that's happening at headquarters, with inconsistent results depending on users' locations.

With the rapid growth of mobile workforces and cloud-based applications, organizations are finding that their remote access VPN is neither optimized for the cloud nor secure. A new approach is necessary to account for today's application mix.

A modern architecture for the mobile workforce

The mobile workforce needs access to the data center and the Internet, as well as applications in the public cloud. A proper architecture should optimize access to all applications, wherever they or your users are located. A SASE solution provides a cloud-delivered networking and security infrastructure that makes it possible for an organization to connect users automatically to a nearby cloud gateway, enable secure access to all applications, and maintain full visibility and inspection of traffic across all ports and protocols (see Figure 2-3).




FIGURE 2-3: Easy access to the connectivity layer, wherever your users are.

For managed devices:

- » Users with managed devices have a SASE client app installed on their laptops, mobile phones, or tablets. The app connects to the SASE platform automatically whenever Internet access is available, without requiring any user interaction.
- » Users can access all their applications, whether in the cloud or the data center. The connectivity layer connects applications in different locations, making it possible to establish secure access (based on application and user identification policies) to public cloud, software as a service (SaaS), and data center applications.
- » SASE delivers protection through the security service layer, such as protections against known and unknown malware, exploits, command-and-control (C2) traffic, and credential-based attacks.

For unmanaged devices:

- » Users with BYOD devices can access applications without an app installed by using a clientless VPN.
- » Clientless VPN also enables secure access to web-based and SaaS applications from unmanaged devices with inline protections by using Security Assertion Markup Language (SAML) proxy integration.

Branch and Retail

Cloud adoption is doing more than changing user mobility strategies. It's affecting branch and retail networking strategies, too. With the growing number of applications in the cloud, it doesn't make sense to carry all of an enterprise network's traffic back to headquarters over expensive multiprotocol label switching (MPLS) connections. As a result, many organizations are adopting new strategies to redesign their wide area networks (WANs) to enable branch offices and retail stores to go directly to the cloud. With the drive to reduce the IT footprint at the branch in order to cut operational costs and reduce complexity, organizations are also looking for ways to reduce the amount of hardware that needs to be physically installed and managed at each location.

The challenges of traditional branch/retail networking

The traditional standard for branch/retail networking uses an MPLS circuit between each remote site and headquarters in a hub-and-spoke topology. This makes sense when the remote site largely uses applications hosted in an internal data center or when bandwidth requirements are not very high. For example, a company that sells machine parts may host an inventory application in its internal data center, and stores across the region might query the database to get real-time information on warehouse inventory. The application does not require significant bandwidth, but the connection must be reliable because any downtime or performance issues could lead to lost business.

Many applications have now moved out of the internal data center and into the public cloud. As a result, hub-and-spoke networking is no longer ideal because traffic passes over the MPLS connection, egresses the perimeter firewall, connects to the cloud-based host, and follows the reverse path back to the user. The MPLS link is a bottleneck because the traffic makes an unnecessary trip to headquarters over a relatively slow connection and adds cost and complexity associated with the additional MPLS resources required to hairpin traffic.

Compounding this issue even further, employees at branch or remote locations need access to more bandwidth-intensive applications than ever before, driving up bandwidth requirements. It's common to see branch offices and retail stores adopt new applications, such as:

- » Employee training applications, often including streaming video and audio content
- » Digital marketing in retail stores to provide product information, personalized shopping tools (for example, gift registries and customization tools), and online catalogs
- » In-store guest Internet

As a result, the transformation to enable direct Internet access at the branch is a necessity for businesses to compete today. However, the options for how it's done can be overwhelming when you consider the need for bandwidth capacity, reliability, operational efficiency, and security.

Augmenting MPLS with direct Internet access

As organizations embark on their cloud journey, traditional connectivity options of private links from branch to data center just don't work. Many organizations augment their private links with Internet connections for direct cloud access. Providing branch locations with direct Internet connections requires IT teams to consider many factors. Plenty of options are available, with most major cities having a range of providers for low-cost, high-speed business-class Internet. However, the top speed of the service is typically not the sole concern. Businesses need to consider the reliability and security of the service as well, and those issues aren't always easy to address.

Several different strategies for augmenting MPLS with direct Internet access can come into play including the following:

- » **MPLS offloading:** MPLS offloading strategies typically don't try to eliminate the MPLS circuit; instead, they reduce the amount of traffic it must carry. For example, many organizations supplement the MPLS circuit with a direct-to-Internet connection so that Internet traffic doesn't have to be backhauled to headquarters.
- » **MPLS replacement with direct to Internet:** In some scenarios, it might make sense to replace the MPLS circuit entirely and implement direct-to-Internet access from the branch. In direct-to-Internet scenarios, the networking requirements can be considerable depending on the topology of the site-to-site VPN connections.
- » **MPLS offloading or replacement with SD-WAN:** Whether you're offloading or replacing MPLS, one of the strategies that you can also adopt is the use of SD-WAN. SD-WAN provides the intelligence to optimize networking decisions based on the applications, networking, and bandwidth requirements that are available, automates complex networking tasks (such as policy-based routing), and provides a common interface to manage networking across branch locations. However, no SD-WAN solution is complete without a natively integrated, robust security system.

A modern architecture for branch transformation

Branch offices need access to all applications, including those in the data center; on the Internet; in SaaS applications; and in public clouds. In other words, the proper architecture should optimize access to all applications, wherever the applications or the users are located.

SASE provides cloud-delivered networking and security infrastructure that makes it possible to connect branch offices to a nearby cloud gateway, enabling secure access to all applications together with full visibility and inspection of traffic across all ports and protocols.

With this architecture, organizations don't have to manage separate on-premises networking and security appliances. Instead of using specialized networking or security hardware, your organization can repurpose an existing branch router or firewall at the branch site, an SD-WAN edge device, or any other Internet Protocol Security (IPSec)-compatible device to connect to the SASE platform. Policies are applied to traffic destined for the cloud, to the Internet, back to corporate headquarters, and even over a full-mesh VPN for branch-to-branch applications.

This immediately eliminates operational expenses, such as the shipping, installation, and ongoing maintenance of extra IT equipment at remote sites. Staffing can focus on operations and protecting the organization from a central location rather than handling the enforcement at the branch network edge.



TIP

In Chapters 3 and 4, you'll learn about the core networking and security capabilities, respectively, that support mobile and remote users and branch/retail use cases in a SASE platform.

IN THIS CHAPTER

- » Defining the need for SD-WANs
- » Getting real about VPNs
- » Implementing ZTNA
- » Ensuring service quality with QoS

Chapter 3

SASE Networking Capabilities

In this chapter, you learn about the core networking capabilities of a SASE solution.

Software-Defined Wide-Area Network

Wide area networks (WANs) use links such as multiprotocol label switching (MPLS), wireless, broadband, virtual private networks (VPNs), and the Internet to give users in remote offices access to corporate applications, services, and resources, allowing them to carry out daily functions regardless of location. Traditional WANs rely on physical routers to connect remote or branch users to applications hosted in data centers. Each router has a data plane (which holds the information) and a control plane (which tells the data where to go). Where the data flows is typically determined by a network engineer or administrator who creates rules and policies, often manually, for each router on the network — a process that can be time-consuming and prone to errors.

Software-defined WAN (SD-WAN) allows enterprises to leverage a wide combination of WAN transport services including MPLS, Long-Term Evolution (LTE), and commodity broadband, to securely connect branches and users to applications both in the cloud and data center.

SD-WAN separates the control and management processes from the underlying networking hardware, making them available as software that can be easily configured and deployed. A centralized control plane means network administrators can create new rules and policies, and then configure and deploy them across an entire network at once.

As cloud applications become mainstream, the traditional approach of a private WAN link backhauling traffic to a data center does not work, because the traffic has to then be sent out to the cloud from the data center. Backhauling traffic to data centers was a suitable WAN architecture when all applications were hosted in data centers. Now that most applications are cloud/software as a service (SaaS) based, it doesn't make sense to backhaul traffic to the data center on its way to the Internet. It's better to go directly to the Internet from the branch (direct Internet access) for cloud/SaaS and back to the data center only for apps hosted there. SD-WAN makes this possible.

Compared to traditional WANs, SD-WANs can intelligently manage multiple types of connections, including MPLS, broadband, Long-Term Evolution (LTE), and others, as well as support applications hosted in data centers, public and private clouds, and SaaS services. SD-WAN can route application traffic over the most optimal path based on performance (that is, latency, jitter, packet loss, and availability), in real time, by intelligently load balancing across multiple links. Prior to SD-WAN, organizations would have to manually configure multiple links to behave a certain way using policy-based routes, for example, to determine what application should take which link.



REMEMBER

Companies are embracing SD-WAN to connect branch offices to the corporate network and provide local Internet breakout for better performance and user experience.



TIP

SD-WAN offers geographically distributed organizations and companies with multiple branches a number of benefits, including

- » **Simplicity:** SD-WAN enables centralized management and simplified configuration rules. In addition, combining SD-WAN with zero-touch provisioning — a feature that helps automate the deployment and configuration processes — organizations can further reduce the complexity, resources, and operating expenses required to turn up new sites.
- » **Greater flexibility and agility:** With SD-WAN, organizations have more connectivity options, such as broadband Internet, which is faster to provision than MPLS. Configuring, deploying, and managing MPLS is time-consuming for most organizations. It can sometimes take a service provider up to three months to install a new MPLS circuit, and MPLS isn't readily available in all areas. SD-WAN remediates this challenge because it separates control of the network services from transport, letting organizations securely use any available Internet connection (such as broadband or LTE) in a given region without being limited to the coverage provided by the MPLS carrier.
- » **Improved user experience:** Without SD-WAN, connecting branch offices to cloud applications is expensive. Traditional WANs must backhaul traffic to the headquarters or corporate data center, usually over MPLS (see Figure 3-1). This can lead to inefficient use of resources and poor performance. By enabling efficient access to cloud-based resources without the need to backhaul traffic to centralized locations, organizations can provide a better overall user experience that leads to less frustration, higher productivity, and better collaboration.
- » **Efficient use of resources:** SD-WAN can lead to greater efficiency in the following ways:
 - *Acquisition of hardware, software, and support:* According to industry research, companies can save up to 40 percent over five years.
 - *Personnel to manage, troubleshoot, and provision WAN equipment.*
 - *Network expenses:* Because SD-WAN supplements or substitutes MPLS with broadband or other Internet connectivity, traffic can be routed based on the best option for cost versus performance.




FIGURE 3-1: Efficient SD-WAN traffic routing.



WARNING

When adopting SD-WAN, however, decision-makers often prioritize connectivity and cost benefits over security. This can put your network at risk.

Although SD-WAN offers many benefits, it also brings many challenges, including new security risks, unreliable performance, and increased complexity resulting from the need for multiple overlays. When security is an afterthought, it tends to be bolted on, introducing management complexity and subpar protection. Moreover, network performance becomes less reliable because organizations use the congested public Internet as the WAN middle mile. Organizations sometimes try to address these challenges by building their own SD-WAN hub and interconnect infrastructures, which results in more complexity.

In a SASE solution, SD-WAN edge devices can be connected to a cloud-based infrastructure, rather than physical SD-WAN hubs located in data center or colocation facilities. This enables the interconnectivity between branch offices without the complexity of deploying and managing physical SD-WAN hubs.



TIP

SD-WAN should be something you're already considering or you've already adopted into your organization's network infrastructure as a way to securely connect and control access to branch offices and remote employees. SASE creates a unified framework for SD-WAN and security services to connect to, providing a single point of view and simplified management solution to protect your network.

Virtual Private Network

A VPN uses a public network, such as the Internet, to enable remote users and sites to connect securely to the corporate network. Two types of VPNs are a remote-access VPN and a site-to-site VPN. Corporate networks are sometimes built on site-to-site VPNs, where the local area network (LAN) of each location can be connected to the data center via a secure WAN on which company resources can be shared. Remote-access VPNs allow individual users to connect to the corporate network remotely.

On VPNs, data travels over the Internet securely through a tunneling protocol, where it's encrypted using Internet Protocol Security (IPSec) or Secure Sockets Layer (SSL). The tunneling protocol also *encapsulates* (wraps) the data with routing information for the receiving user.

Organizations rely on VPNs to provide a secure encrypted connection for mobile users and branch offices to access corporate data, applications, and Internet access. VPNs are not optimized for access to the cloud, however, resulting in no security or access control when users disconnect to reach cloud apps or services (see Figure 3-2).




FIGURE 3-2: Remote-access VPN is not designed to support cloud applications.

A SASE solution encompasses VPN services and enhances the capabilities to operate in a cloud-based infrastructure in order to securely route traffic to public cloud services including SaaS, platform as a service (PaaS), and infrastructure as a service (IaaS), as well as Internet and private cloud apps and services. In an IPsec VPN example, you can create a site-to-site connection to a cloud-based infrastructure from any IPsec-compatible device located at a branch or retail location via a branch router, wireless access point, SD-WAN edge device, or firewall (see Figure 3-3). Mobile users employ an always-on IPsec or SSL VPN connection between their endpoint or mobile device, and a SASE solution ensures consistent traffic encryption and threat prevention.




FIGURE 3-3: SASE uses cloud infrastructure to connect users to both cloud apps and the data center.




REMEMBER

No matter which type of VPN service you use in your organization, a SASE solution provides a unified cloud infrastructure to connect to, instead of backhauling to a VPN gateway at corporate headquarters. This dramatically simplifies the management and policy control needed to enforce least-privilege access rules.

Zero Trust Network Access

As shown in Figure 3-4, companies still lack the necessary security protections and policies needed to adequately protect their users and data. Zero Trust network access (ZTNA) is a key part of the Zero Trust philosophy of “never trust, always verify,” developed by Forrester Research. ZTNA requires users who want to connect to the cloud to authenticate and have their traffic inspected up through Layer 7 via a gateway prior to gaining access to the applications they need. This provides an IT admin with the ability to identify users and create policies to restrict access, minimize data loss, and quickly mitigate any issues or threats that may arise. Many ZTNA products are based on micro-perimeter architectures, which do not provide content inspection, thus creating a discrepancy in the types of protection available for each application. In terms of consistent protection, the organization must build additional controls on top of the ZTNA model and establish inspection for all traffic across all applications.



Source: Verizon Mobile Security Index 2019 report

FIGURE 3-4: Which of the following match your organization's security policies?



REMEMBER

Layer 7 inspection and control are imperative to Zero Trust.

SASE builds upon the key principles of ZTNA and applies them across all the other services within a SASE solution. By identifying users, devices, and applications, no matter where they're connecting from, policy creation and management is simplified. SASE removes the complexity of connecting to a gateway, by incorporating the networking services into a single unified cloud infrastructure.



REMEMBER

A SASE solution should support ZTNA capabilities for protecting applications, as well as apply other security services for the consistent enforcement of data loss prevention (DLP) and threat prevention policies. This is necessary because access controls, in and of themselves, are useful for establishing who the person is, but other security controls are also necessary to make sure their behaviors and actions are not harmful to the organization and its data. It's also necessary to apply the same controls across access to all applications.

WHAT IS ZERO TRUST?

Zero Trust is a cybersecurity strategy that helps prevent successful data breaches by eliminating the concept of trust from an organization's network architecture. Rooted in the principle of "never trust, always verify," Zero Trust is designed to protect modern digital environments by leveraging network segmentation, preventing lateral movement, providing Layer 7 threat prevention, and simplifying granular user-access control.

Zero Trust was created by John Kindervag at Forrester Research, based on the realization that traditional security models operate on the outdated assumption that everything inside an organization's network should be trusted. Under this broken trust model, it's assumed that a user's identity is not compromised and that all users act responsibly and can be trusted. The Zero Trust model recognizes that trust is a vulnerability. Once on the network, users — including threat actors and malicious insiders — are free to move laterally and access or exfiltrate whatever data they aren't limited to. **Remember:** The point of infiltration of an attack is often not the target location.

Quality of Service

As organizations transition from MPLS to SD-WAN using direct Internet access (DIA) links, they're finding that the service quality varies. Quality of service (QoS) establishes bandwidth allocation assigned to particular apps and services. Businesses rely on QoS to ensure that their critical apps and services (for example, medical equipment or credit card processing services) perform adequately. If these systems were to get bogged down due to a lack of available bandwidth caused by network congestion (for example, non-business-related streaming video), this would severely impact business operations and sales (as shown in Figure 3-5).




FIGURE 3-5: Bandwidth without QoS control.



WARNING

Broadband Internet is a “best effort” service that doesn’t provide a consistent bandwidth level. For this reason, QoS isn’t supported on broadband Internet links. If you have QoS configured for your network, your broadband Internet service provider (ISP) will ignore QoS tagging on its routers.

QoS prioritizes business-critical apps, based on a ranking system, so you can choose which apps and services take precedence over others (see Figure 3-6). QoS is an important step when you begin migrating from MPLS. A SASE solution incorporates QoS services in the cloud, allowing you to easily mark sensitive applications, such as voice over IP (VoIP), as high priority over general Internet and entertainment sites and apps.




FIGURE 3-6: Bandwidth with QoS control.



REMEMBER

QoS is immensely important for businesses of any size, especially as they migrate from MPLS to direct to Internet. Managing the QoS traffic and allocation doesn't need to be difficult. A good SASE solution will enable you to dynamically shape traffic based on the policies that prioritize critical application requirements.

IN THIS CHAPTER

- » Deploying a next-generation firewall “as a service”
- » Securing DNS resolution
- » Leveraging threat prevention tools
- » Blocking malicious websites with an SWG
- » Preventing sensitive data loss and ensuring regulatory compliance
- » Identifying and securing access to SaaS apps

Chapter 4

SASE Security Capabilities

In this chapter you learn about the core security capabilities in a SASE solution.

Firewall as a Service

Firewalls were originally designed to protect on-site company networks, but as more companies moved their applications and data to the cloud, firewalls had to evolve. Now, firewall as a service (FWaaS) enables firewalls to be delivered as a cloud service.

In the past, organizations ran all their applications and data in on-site data centers and used a perimeter-based defense to secure their networks, with on-premises firewalls serving as the main security checkpoints. However, as companies moved to the cloud, added more company- and employee-owned mobile devices to their networks, and began using more software as a service (SaaS)

applications and data hosted on third-party infrastructure, they quickly discovered they no longer had clearly defined network perimeters.

They also found that because many of their applications and data were now being run and managed on third-party infrastructure, they no longer had full visibility into, or control over, their entire networks. This problem was further exacerbated by the proliferation of third-party point products that had to be separately managed.

This forced many organizations to completely rethink their approach to security. FWaaS is a deployment method for delivering a firewall as a cloud-based service. FWaaS has the same features of a next-generation firewall, but it's implemented in the cloud. By moving the firewall to the cloud, organizations can benefit from cost savings by eliminating the need to install or maintain security hardware or software across their entire organization.

The FWaaS approach enables organizations to:

- » Aggregate all traffic from multiple sources (for example, on-site data centers, branch offices, mobile users, cloud infrastructure) into the cloud
- » Consistently apply and enforce security policies (fewer error-prone, manual configurations) across all locations and users
- » Gain complete visibility into and control over their networks without having to deploy physical appliances, thereby reducing support costs



TIP

A company with 500 employees can expect to save 37 percent, on average, by using FWaaS solutions versus traditional hardware, according to Secure Data.

A SASE solution incorporates FWaaS into its unified platform, providing the same services as a next-generation firewall but as a cloud-delivered service. By encompassing the FWaaS service model within a SASE framework, organizations can easily manage their deployments from a single platform.




A SASE solution should enable FWaaS capabilities in order to provide the protection of a next-generation firewall by implementing network security policy in the cloud. It's important to ensure your SASE solution doesn't only provide basic port blocking or minimal firewall protections. You need the capabilities of a next-generation firewall, as well as cloud-based security services, such as threat prevention services and Domain Name System (DNS) security.

Domain Name System Security

Each device connected to the Internet has an Internet Protocol (IP) address. The DNS is a protocol that translates a user-friendly domain name, such as `www.paloaltonetworks.com`, to an IP address — in this case, `199.167.52.137`. DNS is ubiquitous across the Internet. Without it, we'd have to memorize random strings of numbers, which our brains aren't equipped to do very well.

DNS is an open service, and by default it doesn't have a way to detect DNS-based threats. As a result, malicious activity within DNS can be used to propagate an attack causing costly damage and downtime (see Figure 4-1).



Source: IDC 2019 Global DNS Threat Report

FIGURE 4-1: DNS attacks are prevalent and result in costly damage and application downtime for organizations.

DNS is a massive and often overlooked attack surface present in every organization. According to the Palo Alto Networks Unit 42 threat research team, almost 80 percent of malware uses DNS to initiate command-and-control (C2) communications (see the “DNS-based attacks: OilRig” sidebar in this chapter). Unfortunately, security teams often lack basic visibility into how threats use DNS to maintain control of infected devices. Adversaries take advantage of the ubiquitous nature of DNS to abuse it at multiple points of an attack, including reliable C2.

Security teams struggle to keep up with new malicious domains and enforce consistent protections for millions of emerging domains at once. It’s impossible for enterprise network and security teams to keep up with the high volume of malicious domains, let alone advanced tactics like DNS tunneling for stealthy data theft.

DNS-BASED ATTACKS: OilRig

OilRig is an active, organized threat group first discovered by the Palo Alto Networks Unit 42 threat research team. Operating primarily in the Middle East, OilRig carefully targets organizations to further its regional strategic goals across multiple industries, including supply-chain-based attacks. As part of its adversary playbook, the group employs sophisticated, custom DNS tunneling for C2 and data exfiltration. The use of tunneling includes

- **ALMA Communicator Trojan**, which uses DNS tunneling to receive commands from the adversary and exfiltrate data. The malware employs specially crafted subdomains to send data to the C2 server and specific Internet Protocol version 4 (IPv4) addresses to transmit data from the C2 to the Trojan over DNS requests.
- **Helminth PowerShell-based Trojan**, which can obtain files from a C2 server using a series of DNS text (TXT) queries repeated every 50 milliseconds, essentially building malware on victim systems through hard-to-detect increments sent over DNS.

OilRig’s use of DNS tunneling allows the group to establish reliable C2 that can potentially evade existing defenses to carry out further stages of the attack.

DNS security protects users by detecting and blocking malicious domains while neutralizing threats. A SASE solution embraces DNS security features by providing consistent security across the network and users, no matter their location, with advanced capabilities that include enabling organizations to:

» **Automatically protect against tens of millions of malicious domains identified with real-time analysis and continuously growing, global threat intelligence:**

Protection continues to grow with data from a large, expanding threat intelligence sharing community. A malicious domain database is created from multiple sources, including the following:

- *Malware prevention* to find new C2 domains, file download source domains, and domains in malicious email links
- *URL filtering* to continuously crawl newfound or uncategorized sites for threat indicators
- *Passive DNS and device telemetry* to understand domain resolution history
- *Threat research* to provide human-driven adversary tracking and malware reverse engineering, including insight from globally deployed honeypots
- *Third-party threat intelligence* sources

» **Predict and stop malicious domains from domain generation algorithm-based malware with instant enforcement.**

Malware's use of domain generation algorithms (DGAs) continues to grow, limiting the effectiveness of blocking known malicious domains alone. DGA malware uses a list of randomly generated domains for C2, which can overwhelm the signature capability of traditional security approaches. DNS security deals with DGA malware by using:

- *Machine learning* to detect new and never-before-seen DGA domains by analyzing DNS queries as they're performed.
- *Easy-to-set policy* for dynamic action to block DGA domains or sinkhole DNS queries.
- *Threat attribution and context* to identify the malware family with machine learning for faster investigation efforts.



Your SASE solution should provide DNS security delivered within the cloud environment as part of the network access. DNS security should be built in, rather than bolted on, to the solution your branch offices and mobile users use to connect to the Internet. The DNS security provided in your SASE solution should leverage a combination of predictive analytics, machine learning, and automation to combat threats in DNS traffic.

Threat Prevention

The dynamic nature of public cloud usage and user mobility requires security teams to adapt and embrace a new approach to threat prevention. According to respondents in a recent ESG survey, threat detection and response is more difficult today than ever before because:

- » The volume and/or sophistication of threats has increased (34 percent).
- » The threat detection/response workload has increased (17 percent).
- » The attack surface has grown (16 percent).
- » Threat detection/response is dependent on many manual processes within the organization (11 percent).
- » The organization uses numerous disparate threat detection/response tools (11 percent).
- » The organization doesn't have the skills or appropriately sized cybersecurity staff (8 percent).

In today's world of small- and large-scale breaches, threat prevention is key to protecting your organization's data and employees. There are a variety of threat prevention tools out there, from anti-malware and intrusion prevention to SSL decryption and file blocking, providing organizations ways to block threats. However, these point products require separate solutions, making management and integration difficult.

Within a SASE solution, all these point products and services are integrated into a single cloud platform. This provides simplified management and oversight of all threats and vulnerabilities across your network and cloud environments.

Stopping exploits and malware by using the latest threat intelligence is crucial to protecting your employees and data. Your SASE solution should incorporate threat prevention tools into its service so you can react quickly and effectively to remediate threats. Be sure to check the quality of threat intelligence that is being provided by the vendor. The vendor should be gathering and sharing data from various sources, including customers, vendors, and other relevant thought leaders, to provide continuous protection from unknown threats.



REMEMBER

Continuous and effective threat prevention, detection, and automated response across your environment requires the following:


- » Granular visibility into your users, apps, and data
- » Advanced threat prevention over the network
- » Threat detection and analysis by correlating risky configurations, anomalous user and network activity, host vulnerabilities, and threat intelligence gathered from multiple data sources
- » Automated response to simplify security event triage
- » Cloud context to expedite security investigations

Secure Web Gateway

In recent years, the emergence of secure web gateways (SWGs) has provided one approach to the problem of securing web traffic from a branch or mobile user's endpoint (discussed in Chapters 2 and 3).

Instead of performing full inspection of all network traffic, a web gateway examines traffic from a web browser and blocks websites and known malware. Organizations looking for a better solution, as opposed to no inspection, may use this approach without having to deploy a hardware appliance at the branch.

Many organizations rely on an SWG to protect employees and devices from accessing malicious websites. According to the *Google Transparency Safe Browsing Report*, more than 500,000 unsafe websites were detected by Google in July 2019 (see Figure 4-2).



■ Malware sites □ Phishing sites

Source: Google Transparency Safe Browsing Report

FIGURE 4-2: More than 500,000 unsafe websites were detected by Google in July 2019.

SWG can be used to block inappropriate content (for example, pornography and gambling) or websites that businesses simply don't want users accessing while at work, such as streaming services like Netflix. Additionally, SWG can be used to enforce an acceptable use policy (AUP) before Internet access is granted.



WARNING

Web gateways are not a substitute for a firewall. Partially inspecting traffic means the remaining traffic passes through uninspected, or else the application breaks. The organization remains blind to applications that legitimately use alternative ports, as well as those intentionally evading inspection. Security is compromised because there is no inspection of non-browser traffic and no protection against other stages of the attack life cycle (see the “Know your enemy: Modern cyberattack strategy” sidebar), such as secondary malware payloads or ongoing C2 traffic with a compromised endpoint.

SWG is just one of the many security services that a SASE solution must provide. As organizations grow and add more and more remote users, coverage and protection becomes more difficult. A SASE solution moves SWG into the cloud, providing protection in the cloud through a unified platform for complete visibility and control over the entire network.



REMEMBER

A SASE solution includes the same security services in an SWG, allowing organizations to control access to the web and enforce security policies that protect users from hostile websites. Remember that SWG is just one service of the SASE solution. Other security services like FWaaS, DNS security, threat prevention, data loss prevention (DLP), and cloud access security broker (CASB) are also necessary.

KNOW YOUR ENEMY: MODERN CYBERATTACK STRATEGY

Modern cyberattack strategy employs a patient, multistep, covert process that blends exploits, malware, and evasion in a coordinated attack. The cyberattack life cycle (see the figure) is a sequence of events that an attacker goes through to successfully infiltrate an organization's network and steal data.



Here are the steps of the cyberattack life cycle:

- 1. Reconnaissance.** Like common criminals, cybercriminals carefully study their victims and plan their attacks, often using social engineering, phishing, email address harvesting, and other tactics to research, identify, and select targets. They also use various tools to scan networks (and SaaS applications) for vulnerabilities, services, and applications that can be exploited.
- 2. Weaponization and delivery.** Next, the attacker determines the malware payload and the method that will be used to deliver it. For example, data files or web pages can be weaponized with exploits that are used to target the victim's vulnerable software and delivered via an email attachment or drive-by-download.
- 3. Exploitation.** The attacker generally has two options for exploitation: social engineering or software exploits. *Social engineering* is a relatively simple technique used to lure someone into clicking on

(continued)

(continued)

a bad link or opening a malicious executable file, for example. *Software exploits* are a more sophisticated technique because they essentially trick the operating system (OS), browser, or other third-party software into running an attacker's code. This means the attacker has to craft an exploit to target specific vulnerable software on the endpoint. When exploitation has succeeded, an advanced malware payload can be installed.

- 4. Installation.** When a target endpoint has been infiltrated, the attacker needs to ensure *persistence* (resilience or survivability). Various types of advanced malware are used for this purpose, including anti-AV software, backdoors, bootkits, and rootkits.
- 5. Command-and-control (C2).** Communication is the life blood of a successful attack. Attackers must be able to communicate with infected systems to enable C2, and to extract stolen data from a target system or network. This communication can also be used by the attacker to move laterally, targeting other systems on the victim's network. C2 communications must be stealthy and can't raise any suspicion on the network.
- 6. Actions on the objective.** Attackers have many different motives for an attack, including data theft, destruction of critical infrastructure, hacktivism, or cyberterrorism. This final phase of the attack often lasts months or even years, particularly when the objective is data theft, because the attacker uses a low-and-slow attack strategy to avoid detection.

Data Loss Prevention

Companies are processing massive amounts of data in more places than ever (for example, in their offices, in the cloud, in multiple SaaS applications and cloud storage environments, and so on). In addition, with cloud and mobile computing technologies, employees now have the ability to directly access applications and data anytime, anywhere, and from any device.

The challenge is:

- » Most companies don't have much visibility into where their sensitive and regulated data is, and how and where their employees access it, use it, or share it with others. In cloud environments, all of the above is especially true.

- » SaaS and public cloud providers may offer some data protection capabilities, which can lead to ineffective and inconsistent security.
- » The number of data breaches by insider threats continues to increase.

To overcome these challenges, it's crucial for companies to put a solid data protection strategy in place.

DLP protects sensitive data (for example, intellectual property, financial data, identities, regulated data, and so on) from loss and theft.



REMEMBER

DLP allows a company to

- » Discover all their sensitive data consistently across different repositories and communication vectors, such as Office 365, Box, Slack, corporate devices, network traffic, and so on.
- » Monitor usage of sensitive data.
- » Protect sensitive data and proactively prevent data leakage.



TIP

For DLP to be effective, companies must

- » Protect their data across their networks, clouds, and users, including SaaS applications, cloud storage, and network traffic.
- » Optimize their DLP deployment and management efforts.
- » Discover, classify, monitor, and protect all their sensitive data, such as personally identifiable information (PII) and intellectual property (IP).
- » Clearly define and enforce policies in order to accurately detect data exposure and violations.
- » Ensure that their data is being stored, accessed, and used in a way that complies with regulations and privacy laws, such as the European Union General Data Protection Regulation (GDPR), California Consumer Privacy Act (CCPA), U.S. Health Insurance Portability and Accountability Act (HIPAA), Payment Card Industry Data Security Standards (PCI DSS), and so on.

DLP is traditionally a composite solution that monitors data within the environments where it's deployed (such as network, endpoints, and cloud). With SASE, DLP becomes a single

cloud-delivered solution centered around the data itself. Policies are consistently applied to sensitive data at rest, in motion, and in use, regardless of its location. With SASE, organizations can finally enable a comprehensive data protection solution that relies on a scalable and simple architecture and allows effective machine learning by leveraging access to all the organization's traffic and data.



REMEMBER

DLP is a necessary tool to protect sensitive data and ensure compliance throughout the organization. With SASE, DLP is an embedded, cloud-delivered service used to accurately and consistently identify, monitor, and protect sensitive data across networks, clouds, and users.

Cloud Access Security Broker

SaaS applications (like Office 365, Box, Slack, and Salesforce) offer companies, employees, and customers many operational benefits.

However, for each positive, there is also a negative when it comes to information security (see Table 4-1).

TABLE 4-1 The Pros and Cons of SaaS Adoption

Pros	Cons
SaaS apps can be deployed quickly. As a software solution, the installation and configuration of SaaS apps are quick and painless. By utilizing the cloud, the apps are easily accessible directly to all users.	Anyone with a credit card can start using almost any cloud service. Cloud services are typically set up without IT and security oversight. Users are able to access the application from anywhere and on any device — secure or not.
Large amounts of data can be stored in the cloud with a low total cost of ownership (TCO) for the organization. Data in the cloud is easily shareable among users within the organization and with external third parties.	Data stored in SaaS applications is practically invisible to IT and can be excessively shared and exposed to more users and threats. When such data is sensitive, it's a huge risk for breaches and noncompliance.

Pros	Cons
They're simple to maintain. Instead of having your IT department manually upgrade the app, that responsibility falls to the SaaS vendors, saving you IT resources.	Maintenance isn't always for the purpose of increasing uptime. SaaS vendors do an amazing job releasing new features and functionality, but this frequent pace of change also makes it difficult for IT and security teams to keep tabs on configurations and risk.
Because SaaS apps live in the cloud, they're scalable, no matter the size of your organization, and remote users can access SaaS apps no matter their location.	Most Tier 1 SaaS apps are designed to be infinitely scalable in theory. The downside is that unsanctioned apps will grow virally in your organization.

Given the ease of use inherent to SaaS apps, the volume and sensitivity of data being transferred, stored, and shared in SaaS cloud environments continues to increase. Simultaneously, users are constantly moving to different physical locations, using multiple devices, operating systems, and application versions to access the data they need.

As a result, some undesirable security tradeoffs have emerged:

- » **Lack of visibility** (and therefore protection) into data uploaded and created in the cloud.
- » **Direct access to applications and data from any device** (including unmanaged devices and bring your own device, or BYOD, personal devices) and from anywhere (including unsafe public Wi-Fi in coffee shops or at home).
- » **Shadow IT**, in which employees use and access unsanctioned applications to get their work done as a workaround to sanctioned IT if they aren't provided with the tools that they need. Shadow IT introduces security risk and potentially exposes sensitive corporate data.



REMEMBER

Shadow IT refers to IT applications and services that are acquired and operated by end users without explicit organizational approval and often without organizational IT knowledge or support.

Many organizations depend on SaaS security, like cloud access security brokers (CASBs), to gain visibility into SaaS application usage (both sanctioned and shadow IT), understand where their

sensitive data resides, enforce company policies for user access, and protect their data from threat actors. CASBs are cloud-based security policy enforcement points that provide a gateway for your SaaS provider and your employees.

However, legacy CASB approaches to securing SaaS applications use a standalone proxy designed to perform a limited amount of inline inspection capabilities. There are different deployment modes by which a CASB can deliver its functions, including network inline, Security Assertion Markup Language (SAML) proxy, agent, and application programming interface (API) based (or introspection). And although CASB can also be used for API-based controls, it often has a limited set of ties to contextual policies about which specific users or devices have access to particular data. Despite multiple options for deployment, there are shortcomings with traditional implementation methods, and many enterprise CASB projects have struggled to get off the ground because of it.



TIP

SaaS security functionality is integrated into SASE as a core capability providing SaaS application and data security in a single platform. A SASE solution helps you understand which SaaS apps are being used and where data is going, no matter where users are located. Specific capabilities should include the following:

»» **SaaS visibility**

- Discovery of shadow IT
- App discovery
- App usage reporting
- App risk assessment
- Configuration assessment

»» **Control and compliance**

- App access control
- Data discovery and classification
- Compliance reporting and remediation
- Unmanaged device access control

»» **SaaS protection**

- Threat protection
- Data protection

- Encryption
- Rights management
- User anomaly detection
- Workflow integration

A SASE solution should incorporate both in-line and API-based SaaS controls for governance, access controls, and data protection. Also called a multi-mode CASB, the combination of in-line and API-based SaaS security capabilities provide superior visibility, management, security, and zero-day protection against emerging threats.

IN THIS CHAPTER

- » Getting full visibility and control of users, data, and apps
- » Simplifying monitoring and reporting
- » Protecting mobile and remote users and enabling consistent security
- » Reducing costs and integration nightmares
- » Improving performance and aligning networking and security

Chapter 5

Ten Benefits of SASE

Here are ten important business and technical benefits of deploying secure access service edge (SASE) in your organization.

Complete Visibility across Hybrid Environments

SASE enables complete visibility of hybrid enterprise network environments that connect data centers, headquarters, branch and retail locations, public and private cloud, and mobile users.

The combination of firewall as a service (FWaaS), secure web gateway (SWG), data loss prevention (DLP), and cloud access security broker (CASB) capabilities and functions in SASE empower enterprise security teams with full visibility into all network activity in the environment, including users, data, and apps.

Control of Users, Data, and Apps

Users are increasingly leveraging a variety of applications — including SaaS applications from multiple devices and locations — for work-related (as well as personal) purposes. Many applications, such as instant messaging (IM), peer-to-peer (P2P) file sharing, and Voice over Internet Protocol (VoIP), are capable of operating on nonstandard ports or hopping ports. Some of these applications are sanctioned by the organization, others are tolerated, and others are unsanctioned. Users are increasingly savvy enough to force applications to run over nonstandard ports through protocols such as Remote Desktop Protocol (RDP) and Secure Shell (SSH), regardless of the organization’s policy regarding various applications (sanctioned, tolerated, unsanctioned).

SASE can classify traffic by application on all ports, by default — and it doesn’t create an administrative burden by requiring you to research which applications use which ports in order to configure appropriate policies and rules. SASE provides complete visibility into application usage along with capabilities to understand and control their use (see Figure 5-1).




FIGURE 5-1: Control application usage in policy.

Monitoring and Reporting

SASE eliminates the need to monitor multiple consoles across different networking and security products and creating separate reports for key metrics. Monitoring and reporting can be done from a “single pane of glass” in SASE, which also helps networking and security teams correlate events and alerts to simplify troubleshooting and accelerate incident response.

Less Complexity

SASE enables your business to simplify networking and security by

- » Eliminating unnecessary, limited use of siloed point security solutions
- » Operating from the cloud to cut operational complexity and cost
- » Avoiding logistical issues with shipping, installing, and upgrading multiple networking and security hardware devices to remote branch (or retail) locations

Consistent Data Protection Everywhere

In a traditional multiprotocol label switching (MPLS) wide area network (WAN), all traffic from branch and retail locations is backhauled to a headquarters or data center location — typically, a headquarters office or on-premises data center. This includes data center and Internet traffic. This basic design architecture eliminates the need for firewalls at branch and retail locations because all traffic can be inspected and a centralized security policy can be enforced by the perimeter firewall at the headquarters or data center location. Of course, this also means that the perimeter firewall can become a bottleneck for the entire enterprise with all traffic flowing through the headquarters or data center.

Consistent data protection is about consolidating data protection policies across every environment and data communication vector. No more disjointed data protection policy and configurations for different SaaS apps, for on-premises repositories, and so on, which cause security blind spots, complex manageability, policy inconsistency, shadow IT, and shadow data. SASE enables a consistent DLP protection policy across every environment where data lives and flows, regardless of its location. New security services and applications with specific security policies can also be rapidly and easily deployed from the cloud to branch and retail locations, instead of having to be individually managed at each location.

Reducing Costs

Organizations may choose to invest in commodity point networking and security products. Although this may initially seem to be a less expensive solution, administrative costs will quickly grow out of control as limited networking and security staff resources must learn different management consoles and operating systems — many of which will potentially have very limited remote management capabilities.

SASE enables organizations to extend the networking and security stack to all their locations in a cost-effective manner via a converged, cloud-delivered solution that fully integrates networking and security capabilities and functions.

Lower Administrative Time and Effort

Managing multiple point networking and security products from different vendors in a large number of locations is an administrative burden that few organizations can afford. The cost to train and retain networking and security staff on a multitude of point networking and security products will quickly exceed the organization's capital investments for these products.

SASE enables single-pane-of-glass management of networking and security functions for all your locations in a consistent manner, which reduces the administrative burden and helps to lower training and retention costs.

Reducing Need for Integration

SASE combines multiple networking and security capabilities and functions in a unified cloud-delivered solution, thereby eliminating the need for complex integrations between multiple point networking and security products from different vendors. See Chapter 1 to learn more about the core networking and security capabilities in SASE.

Better Network Performance and Reliability

SASE helps organizations improve network performance and reliability for all users and locations by delivering SD-WAN capabilities that enable multiple links from different sources — including MPLS, broadband, Long-Term Evolution (LTE), satellite, and more — to be load balanced, aggregated, and or configured for failover. This helps reduce congestion and latency associated with backhauling Internet traffic across MPLS connections or routing traffic across a connection that is experiencing high utilization or performance issues.

Greater Agility

Instead of waiting months for MPLS links to be installed, organizations can rapidly connect branch locations using any available Internet connection — such as broadband or LTE — from a local Internet service provider (ISP) with the networking and security capabilities in SASE.

Glossary

acceptable use policy (AUP): An information security policy that defines appropriate and inappropriate user behavior with respect to content in applications such as web browsing, email, and mobile devices.

Active Directory (AD): A directory service developed by Microsoft for identifying and authenticating users on a Microsoft Windows network or application.

application programming interface (API): A set of protocols, routines, and tools used to develop and integrate applications.

asynchronous transfer mode (ATM): A high-speed, low-latency, packet-switched communications protocol.

bring your own device (BYOD): A mobile device policy that permits employees to use their personal mobile devices in the workplace for work-related and personal business.

California Consumer Privacy Act (CCPA): A privacy rights and consumer protection statute for residents of California that was enacted in 2018 and became effective on January 1, 2020.

cloud access security broker (CASB): Software that monitors activity and enforces security policies on traffic between an organization's users and cloud-based applications and services.

command-and-control (C2): Communications traffic between malware and/or compromised systems and an attacker's remote server infrastructure used to send and receive malicious commands or exfiltrate data.

data loss prevention (DLP): A data protection strategy to detect the unauthorized storage or transmission of sensitive data.

direct Internet access (DIA): A networking strategy to provide broadband Internet access to a remote site. Direct to Internet complements or replaces conventional MPLS hub and spoke topologies. *See also* multiprotocol label switching (MPLS).

DNS hijacking: An attack technique that incorrectly resolves DNS queries to redirect victims to malicious sites. Also known as DNS redirection. *See also* Domain Name System (DNS).

DNS resolver: A server that relays requests for IP addresses to root and top-level domain servers. *See also* DNS root server, top-level domain (TLD), *and* Domain Name System (DNS).

DNS root server: An authoritative name server for a specific TLD in the DNS root zone of the Internet. *See also* top-level domain (TLD) *and* Domain Name System (DNS).

DNS tunneling: An attack technique that exploits the DNS protocol to tunnel malware and other data through a network. *See also* Domain Name System (DNS).

domain generation algorithm (DGA): A program developed by attackers that generates semi-random domain names so that malware can quickly generate a list of domains that it can use for C2 communications. *See also* command-and-control (C2).

Domain Name System (DNS): A hierarchical, decentralized directory service database that converts domain names to IP addresses for computers, services, and other computing resources connected to a network or the Internet.

exploit: Software or code that takes advantage of a vulnerability in an operating system or application, and causes unintended behavior in the operating system or application, such as privilege escalation, remote control, or a denial of service.

Extensible Markup Language (XML): A human- and machine-readable markup language.

firewall as a service (FWaaS): A firewall platform provided as a service offering in a cloud environment.

General Data Protection Regulation (GDPR): A European Union law on data protection and privacy for all individuals within the EU and the European Economic Area. The GDPR supersedes the Data Protection Directive (95/46/EC) and became enforceable in 2018.

Health Insurance Portability and Accountability Act (HIPAA):

U.S. legislation passed in 1996 that, among other things, protects the confidentiality and privacy of protected health information (PHI). *See also* protected health information (PHI).

hybrid cloud: An environment that combines a private cloud (internal data center) with resources in the public cloud. *See also* private cloud *and* public cloud.

infrastructure as a service (IaaS): A category of cloud computing services in which the customer manages operating systems, applications, compute, storage, and networking, but the underlying physical cloud infrastructure is maintained by the service provider.

instant messaging (IM): A type of real-time online chat over the Internet.

intellectual property (IP): Includes patents, trademarks, copyrights, and trade secrets.

Internet Engineering Task Force (IETF): An international, membership-based, nonprofit organization that develops and promotes voluntary Internet standards.

Internet Protocol (IP): The OSI Layer 3 protocol that's the basis of the modern Internet. *See also* Open Systems Interconnection (OSI) model.

Internet Protocol Security (IPSec): An IETF open-standard VPN protocol for secure communications over IP-based public and private networks. *See also* Internet Engineering Task Force *and* virtual private network.

Internet service provider (ISP): A telecommunications company that provides access to the Internet.

intrusion prevention system (IPS): A hardware or software application that both detects and blocks exploits and malicious activity such as C2 traffic. *See also* command-and-control (C2).

Lightweight Directory Access Protocol (LDAP): A protocol for accessing directory services, typically used to identify and authenticate users.

local area network (LAN): A computer network that connects computers in a relatively small area, such as an office building, warehouse, or residence.

Long-Term Evolution (LTE): A type of 4G cellular connection that provides fast connectivity primarily for mobile Internet use.

malware: Malicious software or code that typically damages or disables, takes control of, or steals information from a computer system.

man-in-the-middle attack: A type of attack where the attacker impersonates a legitimate service in order to intercept and sometimes modify communications.

mobile device management (MDM): Software used to manage the administration of mobile devices such as smartphones and tablets.

multi-cloud: An environment that consists of multiple types of clouds (such as a public and private cloud, more commonly known as a hybrid cloud), or multiple vendors of the same type of cloud (such as using Amazon Web Services, Google, and Microsoft for public cloud applications). *See also* hybrid cloud, private cloud, *and* public cloud.

multiprotocol label switching (MPLS): A method of forwarding packets through a network by using labels inserted between Layer 2 and Layer 3 headers in the packet.

Open Systems Interconnection (OSI) model: The seven-layer reference model for networks. The layers are Physical, Data Link, Network, Transport, Session, Presentation, and Application.

Payment Card Industry Data Security Standard (PCI DSS): A proprietary information security standard mandated for organizations that handle American Express, Discover, JCB, MasterCard, or Visa payment cards.

peer-to-peer (P2P): A distributed application architecture that enables sharing between nodes.

Personal Information Protection and Electronic Documents Act (PIPEDA): A Canadian data privacy law that governs how private-sector organizations collect, use, and disclose personal information in the course of conducting commercial business.

personally identifiable information (PII): Data (such as name, address, Social Security number, birthdate, place of employment, and so on) that can be used on its own or with other information to identify, contact, or locate a person.

phishing: A social engineering cyberattack technique widely used in identity theft crimes. An email, purportedly from a known legitimate business (typically financial institutions, online auction sites, retail stores, and so on), requests the recipient to verify personal information online at a forged or hijacked website.

platform as a service (PaaS): A category of cloud computing services in which the customer is provided access to a platform for deploying applications and can manage limited configuration settings, but the operating system, compute, storage, networking, and underlying physical cloud infrastructure is maintained by the service provider.

private cloud: A cloud computing deployment model that consists of a cloud infrastructure that is used exclusively by a single organization.

protected health information (PHI): Any information about health status, healthcare, or healthcare payments that can be associated with a specific, identifiable individual.

public cloud: A cloud computing deployment model that consists of a cloud infrastructure that is open to use by the general public.

quality of service (QoS): The ability to prioritize traffic based on operational needs and importance.

Remote Authentication Dial-in User Service (RADIUS): An open-source, UDP-based client-server protocol used to authenticate remote users. *See also* User Datagram Protocol (UDP).

Remote Desktop Protocol (RDP): A proprietary Microsoft protocol that provides remote access to a computer. RDP uses TCP port 3389 and UDP port 3389 by default. *See also* Transmission Control Protocol (TCP) *and* User Datagram Protocol (UDP).

secure access service edge (SASE): Defined by Gartner as “an emerging offering combining comprehensive WAN capabilities with comprehensive network security functions (such as SWG, CASB, FwaaS, and ZTNA) to support the dynamic secure access needs of digital enterprises”. *See also* wide-area network (WAN), secure web gateway (SWG), cloud access security broker (CASB), firewall as a service (FWaaS), *and* Zero Trust network access (ZTNA).

Secure Shell (SSH): A cryptographic network protocol that provides secure access to a remote computer.

Secure Sockets Layer (SSL): A transport layer protocol that provides session-based encryption and authentication for secure communication between clients and servers on the Internet.

secure web gateway (SWG): A security platform or service that is designed to maintain visibility in web traffic. Additional functionality may include web content filtering.

Security Assertion Markup Language (SAML): An XML-based, open-standard data format for exchanging authentication and authorization credentials between organizations. *See also* Extensible Markup Language (XML).

service-level agreement (SLA): Formal minimum performance standards for systems, applications, networks, or services.

shadow IT: IT applications and services that are acquired by end users without explicit organizational approval and often without organizational IT knowledge or support.

social engineering: A technique for hacking that uses deception to trick the victim into performing an action or revealing sensitive information.

software as a service (SaaS): A category of cloud computing services in which the customer is provided access to a hosted application that is maintained by the service provider.

software-defined perimeter (SDP): A software-defined perimeter extends access to private applications (either in the data center or the public cloud).

software-defined wide-area network (SD-WAN): A newer approach to wide area networking that separates the network control and management processes from the underlying hardware, and makes them available as software.

top-level domain (TLD): A domain at the highest (root) level of the DNS of the Internet. Some examples include .com, .edu, .gov, .net, .org, as well as country code TLDs such as .us and .ca.

Transmission Control Protocol (TCP): A connection-oriented protocol responsible for establishing a connection between two hosts and guaranteeing the delivery of data and packets in the correct order.

unified threat management (UTM): A security appliance that combines a number of services such as firewall, anti-malware, and intrusion prevention capabilities into a single platform.

Uniform Resource Locator (URL): Commonly known as a *web address*. The unique identifier for any resource connected to the web.

User Datagram Protocol (UDP): A network protocol that doesn't guarantee packet delivery or the order of packet delivery over a network.

virtual private network (VPN): An encrypted tunnel that extends a private network over a public network (such as the Internet).

Voice over Internet Protocol (VoIP): Telephony protocols that are designed to transport voice communications over TCP/IP networks.

vulnerability: A bug or flaw in software that creates a security risk that may be exploited by an attacker.

wide area network (WAN): A computer network that spans a wide geographical area and may connect multiple local area networks. *See also* local area network (LAN).

Zero Trust network access (ZTNA): A “never trust, always verify” security approach that ensures proper user context through authentication and attribute verification before allowing access to apps and data in the cloud or data center.



Consistent security, everywhere that you need it

You shouldn't have to compromise between speed and security.

Prisma Access by Palo Alto Networks provides the industry's most comprehensive Secure Access Service Edge (SASE), enabling you to securely embrace cloud and mobility.

Learn how Prisma Access provides connectivity and consistent security to mobile users, branch offices and retail locations, with a personal online demonstration.

<https://www.paloaltonetworks.com/company/request-demo>



Build your digital transformation on a SASE framework

With digital transformation initiatives driving cloud adoption, mobility, and software-defined wide area networks (SD-WANs), the ability for organizations to remain secure and prevent data breaches is becoming difficult. This is especially true when considering that traditional security solutions weren't designed with the cloud in mind, which creates problems with complexity, administrative effort, and incomplete protection. A secure access service edge (SASE) provides connectivity and consistent security to mobile users, branch offices, and retail locations, anywhere in the world.

Inside...

- Learn what a SASE solution is
- Discover ways to reduce costs and integration nightmares
- Understand how to gain full visibility and control over users, apps, and data
- Identify ways to improve performance and align networking and security
- See how a SASE solution can simplify monitoring and reporting



Lawrence Miller served as a Chief Petty Officer in the U.S. Navy and has worked in information technology in various industries for more than 25 years. He is the co-author of *CISSP For Dummies* and has written more than 150 *For Dummies* books on numerous technology and security topics.

Go to **Dummies.com**[™]
for videos, step-by-step photos,
how-to articles, or to shop!

ISBN: 978-1-119-69602-5

Not for resale



for
dummies[®]
A Wiley Brand

WILEY END USER LICENSE AGREEMENT

Go to www.wiley.com/go/eula to access Wiley's ebook EULA.