

LEARNING MADE EASY

Palo Alto Networks 2nd Special Edition

Secure Access Service Edge (SASE)

for
dummies[®]
A Wiley Brand



Reduce networking
and security complexity

—
Stop cyberattacks with
consistent security

—
Deliver exceptional
user experiences

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)

**for
dummies[®]**
A Wiley Brand



Secure Access Service Edge (SASE)

Palo Alto Networks 2nd Special Edition

by Lawrence Miller

**for
dummies**[®]
A Wiley Brand

Secure Access Service Edge (SASE) For Dummies®, Palo Alto Networks 2nd Special Edition

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

Copyright © 2022 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: WHILE THE PUBLISHER AND AUTHORS HAVE USED THEIR BEST EFFORTS IN PREPARING THIS WORK, THEY 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 ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. NO WARRANTY MAY BE CREATED OR EXTENDED BY SALES REPRESENTATIVES, WRITTEN SALES MATERIALS OR PROMOTIONAL STATEMENTS FOR THIS WORK. THE FACT THAT AN ORGANIZATION, WEBSITE, OR PRODUCT IS REFERRED TO IN THIS WORK AS A CITATION AND/OR POTENTIAL SOURCE OF FURTHER INFORMATION DOES NOT MEAN THAT THE PUBLISHER AND AUTHORS ENDORSE THE INFORMATION OR SERVICES THE ORGANIZATION, WEBSITE, OR PRODUCT MAY PROVIDE OR RECOMMENDATIONS IT MAY MAKE. THIS WORK IS SOLD WITH THE UNDERSTANDING THAT THE PUBLISHER IS NOT ENGAGED IN RENDERING PROFESSIONAL SERVICES. THE ADVICE AND STRATEGIES CONTAINED HEREIN MAY NOT BE SUITABLE FOR YOUR SITUATION. YOU SHOULD CONSULT WITH A SPECIALIST WHERE APPROPRIATE. FURTHER, READERS SHOULD BE AWARE THAT WEBSITES LISTED IN THIS WORK MAY HAVE CHANGED OR DISAPPEARED BETWEEN WHEN THIS WORK WAS WRITTEN AND WHEN IT IS READ. NEITHER THE PUBLISHER NOR AUTHORS SHALL BE LIABLE FOR ANY LOSS OF PROFIT OR ANY OTHER COMMERCIAL DAMAGES, INCLUDING BUT NOT LIMITED TO SPECIAL, INCIDENTAL, CONSEQUENTIAL, OR OTHER DAMAGES.

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 978-1-119-89742-2 (pbk); ISBN 978-1-119-89743-9 (ebk)

Publisher's Acknowledgments

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

Project Editor: Elizabeth Kuball

Acquisitions Editor: Ashley Coffey

Editorial Manager: Rev Mengle

Business Development

Representative: Cynthia Tweed

Production Editor:

Saikarthick Kumarasamy

Special Help: Shannon Bonfiglio,

Carmine Clementelli,

Matt De Vincentis, Don Meyer,

Brian Rimmel, Kalie Radsmikham,

Neelum Khan, Rajesh Kari,

Faithe Wempen

Introduction

If you're like most people, the environment you work in today is very different than it was just a few years ago. The rapid shift to cloud and adoption of hybrid work, accelerated by the COVID-19 pandemic, represent two of the most dramatic changes the world has experienced in recent history. We've now entered an era in which the workplace is as likely to be a kitchen table as an office cubicle, and cloud applications are vital to people's productivity. However, although work habits and apps have completely transformed, networks haven't fundamentally changed in more than 20 years.

With increasing numbers of remote users, branch offices, data, and services located outside the traditional corporate network, organizations are struggling to ensure sufficient levels of security and connectivity.

Most network and network security products on the market today weren't designed to handle all the types of traffic and security threats that organizations must deal with now. This forces organizations to adopt multiple products to handle different requirements, such as secure web gateways (SWG), firewalls, virtual private network (VPN) remote access, Multiprotocol Label Switching (MPLS) 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, and a 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 applications — including public cloud, private cloud, software as a service (SaaS), platform as a service (PaaS), and infrastructure as a service (IaaS) — delivered through a common framework to users in business offices, home offices, and remote locations. 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 *Secure Access Service Edge (SASE) For Dummies*, I fill you in on this 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 six chapters that explore the following:

- » Modern trends and their impact on the evolution of networking architectures (Chapter 1)
- » SASE use cases (Chapter 2)
- » SASE networking capabilities (Chapter 3)
- » SASE security capabilities (Chapter 4)
- » Digital experience monitoring (Chapter 5)
- » Ten benefits of SASE (Chapter 6)

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's looking for a better way to simplify your approach to networking and security. 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 80 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/sase.

IN THIS CHAPTER

- » Considering how the pandemic has driven the move to hybrid work
- » Understanding the role of the cloud in digital transformation strategies
- » Evolving the network architecture
- » Discovering a new approach to enterprise networking and security

Chapter 1

The Evolution of Networking

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

Understanding the Pandemic's Effect on Hybrid Work

The pandemic has affected work life as we know it. According to Palo Alto Networks's *State of Hybrid Workforce Security 2021 Report*, 76 percent of global workers want the option to continue working remotely at least part of the time.

This increase in remote workers has also increased the number of devices being used for business. People are using smartphones to access the Internet and software as a service (SaaS) apps, not only for personal computing needs, but also for work.

THE TOP HYBRID WORK SECURITY THREATS



REMEMBER

The hybrid workforce is here to stay, and with that comes security risks that organizations need to consider. Cybercriminals often target the weakest link in a network, so it's important to understand the risks when adapting a hybrid workforce for your organization.

Here are some of the top hybrid work security threats today:

- **Home networks:** Home networks can be insecure and often lack the security capabilities that corporate or branch offices utilize.
- **Phishing:** Phishing attacks still target people by email, and with more people working remotely, there are more opportunities. Also, many people use personal devices for work, which provides a greater opportunity for phishing to take place through mobile channels like text messaging and phony websites.
- **Shadow IT:** As more people work from home or remotely, the use of *shadow IT* (unsanctioned apps) can increase, causing gaps in security that IT departments are unaware of. Applications not managed by an organization's IT (like Google Drive, Slack, and WhatsApp) can open the door to threats.
- **Lost devices:** As more people work remotely, their use of personal devices for work purposes increases as well, including mobile devices, tablets, and laptops. Those devices now have corporate data on them, so if they get lost, the organizations are at risk of losing corporate data.
- **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 are often unintentional, such as when someone inside a company accidentally transfers confidential or sensitive data to an unsanctioned/unapproved cloud application, or when they overshare confidential or sensitive data on cloud sharing apps or public cloud storage. However, intentional leaks also happen, like when an attacker or a disgruntled employee deliberately steals the company's data.

Work is also constantly closer at hand than ever before. The prevalence of public Wi-Fi hotspots and 5G cellular connectivity means that the Internet and work assets are always just a few clicks away. This ubiquitous connectivity enables users to work on their laptops, tablets, and smartphones from anywhere.

Business leaders are increasingly viewing work as an activity, rather than a place, and adopting policies for hybrid work and bring your own device (BYOD) that enable employees to take advantage of these new realities. As the pandemic has shown, 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.

This “new normal” brings with it some new challenges, though. Hybrid work introduces new networking and security concerns that traditional remote access connectivity is not designed to address.

Journeying to the Cloud

We live in an age of cloud and digital transformation. Users and applications have moved outside the traditional network perimeter, accessing an ever-increasing number of applications, including 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 *2021 State of the Cloud Report* from Flexera found that companies are fully embracing multi-cloud, with 92 percent having a multi-cloud strategy and 80 percent having a hybrid cloud strategy. Diving further into the type of hybrid cloud strategy, the most common instance is 43 percent of enterprises using a combination of multiple public and private clouds, while 76 percent incorporate multiple public clouds and 56 percent use more than one private cloud.



REMEMBER

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

Assessing 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 wide area network (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 the perimeter firewall could enforce a centralized security policy.

As Internet usage increased, many branch offices began to experience performance issues and latency. Backhauling their Internet traffic across the MPLS connection for perimeter firewall inspection created significant bottlenecks. This 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 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 it introduced a whole new set

of challenges. On the networking side, these challenges have included:

- » **Routing complexity:** Routers must 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). The simplest solution in most cases 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 while 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 ISP provides a commodity router for the DIA link and doesn't 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 have included:

- » **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 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 multiple vendors mean different management interfaces, different operating systems, and different remote administration tools — multiplied by the number of 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 daily from these different tools.

Introducing the SASE Vision

A new architecture was needed to address this shift in networking and security requirements. Gartner writes about a model known as *secure access service edge*.



REMEMBER

SASE converges networking and security services into one unified, cloud-delivered solution that includes the following, as summarized in Figure 1-1:

- » Networking
 - Software-defined wide area network (SD-WAN)
 - Virtual private networks (VPNs)
 - Quality of service (QoS)
 - Routing
 - SaaS acceleration
- » Security
 - Zero Trust network access (ZTNA)
 - Cloud secure web gateway (Cloud SWG)
 - Cloud access security broker (CASB)
 - Firewall as a service (FWaaS)

- DLP
- Domain Name System (DNS) security
- Threat prevention

» User experience

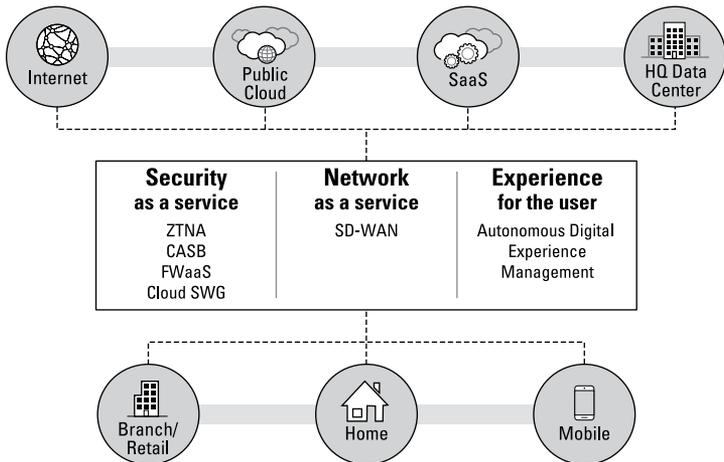


FIGURE 1-1: SASE delivers advanced network and security capabilities in a converged cloud-delivered solution.

Revisiting Modern Networking and Security Challenges with SASE

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

- » **Lower capital costs:** SASE requires 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 users' devices.
- » **Full visibility and control:** SASE provides full visibility and control with cloud-delivered capabilities including ZTNA, cloud SWG, CASB, and FWaaS.

» **Less complexity:** All cloud service management functions 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.



WARNING

Converging networking and security in the cloud with SASE promises to remedy the shortcomings of legacy security and networking architectures. However, many solutions on the market today are incomplete, requiring organizations to make trade-offs between security and functionality or require other products to fill the gaps. Dubbed *multivendor SASE*, this approach retains the legacy challenges of stitching together a multivendor environment, and troubleshooting can be a nightmare. What's the solution? Keep reading to find out.

- » Enabling hybrid workforces with SASE
- » Connecting and securing branch and retail locations

Chapter 2

SASE Use Cases

In this chapter, you find out about some of the most common use cases today for a secure access service edge (SASE), including hybrid workforces and branch locations. Both use cases present different challenges that a SASE solution can solve.

Discovering How SASE Enables Hybrid Workforces

Connecting and securing the hybrid workforce with traditional solutions can be a challenge, especially when users work from home or in locations where you don't have IT staff. For years, the standard solution for connecting remote users into a corporate network was remote-access virtual private networks (VPNs). In fact, for many people, *remote access* and *VPN* are synonymous.

However, the requirements for remote access today are very different than when VPN was invented in the mid-1990s. Today, IT is asked to support the needs of a dynamic, mobile workforce accessing applications that may be hosted in public cloud, private cloud, software as a service (SaaS), or conventional data centers, while maintaining high levels of performance and robust, consistent security controls. This requires an entirely new approach to remote access.

Traditional remote-access VPN limitations

Remote-access VPNs are 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 encrypted 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 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.

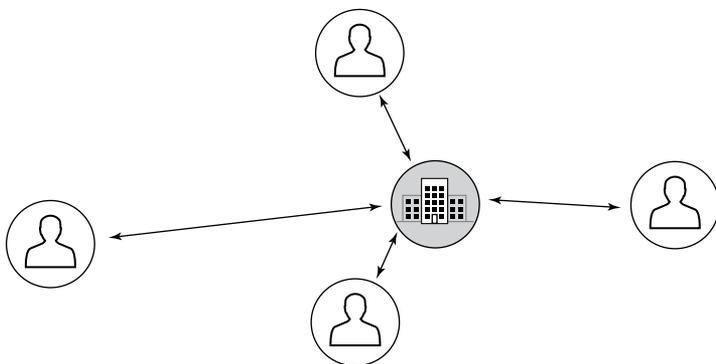


FIGURE 2-1: Traditional remote-access VPN architecture.

The model breaks down when a mixture of cloud applications is involved. With remote-access VPN, traffic always goes to the VPN concentrator or gateway first, even if the application is hosted in the cloud (as shown in 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!) round trip to reach an Internet-accessible location. This is sensible from a security perspective, but it doesn’t make sense for network optimization.

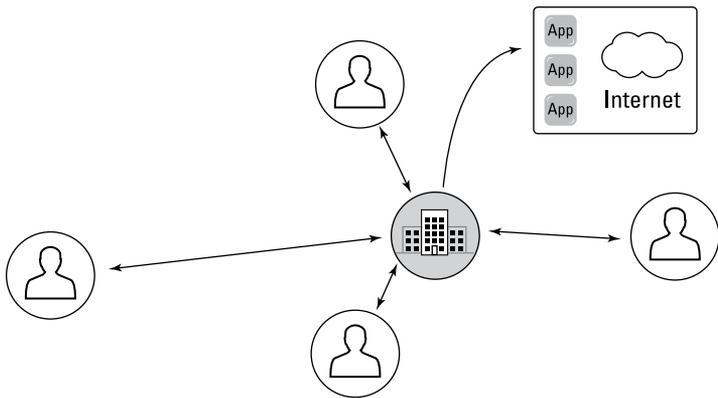


FIGURE 2-2: Traditional remote-access VPN backhauling traffic to reach the cloud.

Using cloud applications over remote-access VPN can hurt the user experience. 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, and that leads to multiple issues.



WARNING

When users are disconnected, their organizations lose visibility into application usage, control over access to unsanctioned applications, and the ability to enforce security policies. In addition, the drastic increase in mobile workforce places significant demands on VPN gateways/concentrators to scale without the infrastructure to support it. The constant traffic overloads force the VPN gateways to deliver poor performance and negatively impact the end-user experience.

Unsatisfactory compromises

To compensate for the networking problems with remote-access VPN, IT teams typically introduce multiple compromises, each with its own security implications:

- » **User-initiated tunnel:** A common remote-access VPN deployment model is to let users initiate the tunnel as needed. They typically 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: Internet and cloud traffic are not inspected.
- » **Web proxy/secure web gateway (SWG):** To compensate for scenarios in which 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 secure nor optimized for the cloud. A new approach is necessary to account for today's application mix.

A modern architecture for the hybrid workforce

Today's hybrid workforce needs access to the data center and the Internet, as well as to applications in the public cloud. A proper architecture should optimize access to all applications, wherever they (and the users) are located.



REMEMBER

A SASE solution provides a cloud-delivered networking and security infrastructure that enables an organization to connect users automatically to a nearby cloud gateway, provide secure access to all applications, and maintain full visibility and inspection of traffic across all ports and protocols (see Figure 2-3).

The benefits are significant for both managed and unmanaged devices.

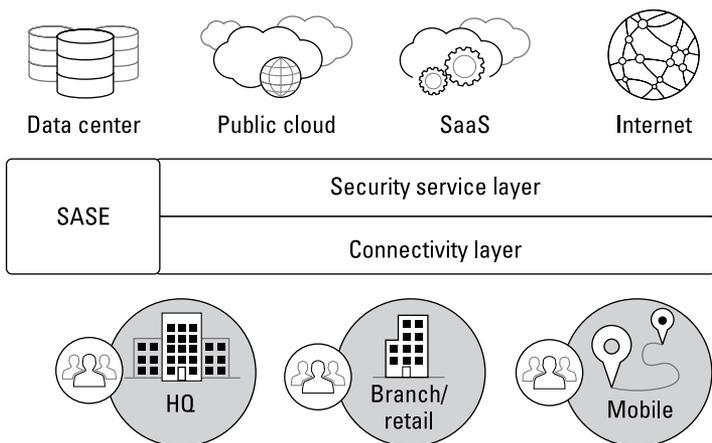


FIGURE 2-3: Easy access to the connectivity layer, wherever your users are.

For managed devices:

- » Users 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, 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.
- » Organizations migrating from a legacy proxy-based web security solution to SASE should have the opportunity to do so without significant network architecture changes. Over time, customers can easily transition from a proxy-based architecture to a more secure connection method that protects all apps, ports, and protocols, not just web.

For unmanaged devices:

- » Users taking advantage of “bring your own device” (BYOD) policies can securely access applications without an app installed by using a clientless VPN.
- » Clientless VPN enables secure access to web-based and SaaS applications from unmanaged devices with inline protections by using Security Assertion Markup Language (SAML) proxy integration.

CASE STUDY: AN ENERGY SERVICES PROVIDER

Companies around the world have adapted to support a hybrid workforce, sometimes overnight. A SASE solution can help, as it did for an energy services provider with more than 100,000 employees in 120 countries looking to support its remote workforce during and after the pandemic.

It needed:

- A solution to scale with its remote workforce quickly
- To stay ahead of cybersecurity risks and threats
- A faster way to connect remote employees and cloud applications

A SASE solution helped the company to:

- Increase the number of employees working remotely from 25,000 to 80,000 users
- Reduce complexity
- Strengthen security and enable a single unified policy base across the enterprise
- Provide a seamless and consistent experience for all users

Enabling Efficient and Secure Branch and Retail Connectivity

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 redesigning 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 to cut operational costs and reduce complexity, organizations are also looking for ways to reduce the amount of hardware that needs to be physically deployed and managed at each location.

The challenges of traditional branch and retail networking

The traditional standard for branch and retail networking uses an MPLS circuit between each remote site and headquarters or the data center 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. Retail stores across the region may 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.



WARNING

Many applications have now moved out of the internal data center and into the public cloud. As a result, hub-and-spoke networking creates serious performance issues because traffic must pass over the MPLS connection, egress the perimeter firewall, connect to the cloud-based host, and then follow 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. This adds cost and complexity due to 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:

- » Real-time collaboration tools such as videoconferencing, instant messaging, file sharing, and Voice over Internet Protocol (VoIP)
- » Video streaming, cloud application access, and online data backup services
- » In-store guest Wi-Fi

As a result, enabling direct Internet access at the branch is necessary 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 have embraced the cloud, traditional connectivity options of private links from branch to data center have begun to create problems. Many organizations have augmented their private links with Internet connections to improve WAN availability and enable direct cloud access.



REMEMBER

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 speed of the service is not the only concern. Organizations also need to consider the reliability and security of the service, and those issues aren't always easy to address.

As a result, many organizations look to software-defined wide area network (SD-WAN) as the answer to these challenges. SD-WAN provides the intelligence to:

- » Optimize forwarding decisions based on applications, transports, bandwidth availability, and performance service-level agreements (SLAs)

- » Automate complex networking tasks (such as policy-based routing)
- » Deploy and configure at scale
- » Provide a centralized interface to manage networking across branch locations

However, no SD-WAN solution is complete without a natively integrated, robust security service.

Replacing MPLS with broadband and 5G



WARNING

As organizations expand, their branch offices become distributed, with more remote locations added as part of their infrastructure. Providing WAN connectivity with MPLS to these branches comes with operational challenges and significant costs.

Similarly, organizations that grow due to mergers and acquisitions create a heterogeneous WAN network. Managing the differences in providers, SLAs, and bandwidth requirements monopolizes IT resources; the results are poor network connectivity and degraded application experience. They require reliable, highly available WAN connectivity that can easily support the bandwidth demands of the cloud applications.

Many organizations are easily transforming complex networks with affordable and high-bandwidth Internet connections like broadband and, most recently, 4G LTE and 5G. With the speed and reliability enhancements and cost savings they provide, metered links like 4G and 5G are proving just as effective as WAN links for remote and mobile locations. Organizations now are able to replace their MPLS WAN connectivity without any compromises in speed or performance.



REMEMBER

Many organizations have found a de facto solution in SD-WAN due to its ability to support carrier-independent multiple WAN links like broadband, direct Internet, and 4G/5G. SD-WAN's automated VPN connectivity delivers encrypted WAN links on top of public Internet connections that ensure security and conformance. In addition, the application intelligence-based steering ensures that the WAN links are best utilized as active-active or active-backup based on bandwidth, performance SLAs, and business policies to deliver the best user experience.

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. The proper architecture should optimize access to all applications, wherever the applications or the users are located.

This architecture is known as the *thin-branch approach*. Much of the branch services are done in the cloud, which keeps the branch lightweight. A thin-branch approach enables businesses to manage security and access through a centralized control via the cloud.

In contrast, legacy WAN solutions often take a *thick-branch approach*, in which much of the branch services such as security, segmentation, routing, and more are done specifically at the branch, requiring more overhead and effort.



WARNING

Using the thick-branch/legacy WAN approach can result in infrastructure sprawl, separate management interfaces, and tedious troubleshooting that can increase operational complexity significantly. In addition, managing WAN connectivity, VPN tunnels, quality of service (QoS), and security policies at the branch demands higher processing power and resources. Businesses are forced to upgrade their branch infrastructures, adding costs to improve application performance and user experience.



REMEMBER

The thin-branch approach aligns heavily with the SASE architecture, utilizing the cloud and providing a positive user experience. 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. 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 change 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 instead of handling the enforcement at the branch network edge.

CASE STUDY: A HIGH-TECH COMPANY

Organizations with branch and retail locations often struggle to provide adequate connectivity and security outside the corporate headquarters or data centers. A SASE solution with SD-WAN can help, as it did for a high-tech company with more than 60,000 employees that was looking to reduce costs and increase network speeds.

The challenges it was facing included:

- An inability to scale or meet employee needs with its legacy MPLS solution
- Unreliable Internet connectivity that impacted branch operations
- Extensive manual operations that consumed IT staff's time and resources

Implementing a SASE solution with SD-WAN provided:

- Application awareness and insights into application traffic
- Centralized management for simplified network operations
- A zone-based firewall for branch segmentation and security
- Improved uptime and availability of branch locations
- Scalability up to 2 Gbps of WAN throughput at large office locations

IN THIS CHAPTER

- » Defining the need for SD-WANs
- » Getting real about VPNs
- » Ensuring service quality with QoS
- » Implementing intelligent routing
- » Accelerating SaaS

Chapter 3

SASE Networking Capabilities

In this chapter, you find out about the core networking capabilities of a secure access service edge (SASE) solution and how security has become a critical component to software-defined wide area network (SD-WAN). In addition, you look at how to ensure the best quality for your SD-WAN.

Discovering How SD-WANs Provide Value

Wide area networks (WANs) use links such as Multiprotocol Label Switching (MPLS), wireless, broadband, virtual private networks (VPNs), and direct Internet to give users in remote offices access to applications, services, and resources, enabling 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. Access rules, traffic policies, and quality of service (QoS) prioritization need to be manually configured on each device. The data flows are typically determined by a network engineer or administrator who creates rules and policies, often manually, for each router on the network. This process can be time-consuming and prone to errors.

SD-WAN enables enterprises to leverage a combination of WAN transport services including MPLS, Long-Term Evolution (LTE), 5G, and commodity broadband to securely connect branches and users to applications both in the cloud and in the data center.

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



REMEMBER

As cloud applications become mainstream, the traditional approach of a private WAN link backhauling traffic to a data center doesn't work, because the traffic must 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. However, 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, or DIA) 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, 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 (considering factors such as latency, jitter, packet loss, availability, and more), in real time, by intelligently load-balancing across multiple links. Prior to SD-WAN, organizations had to manually configure multiple links to behave a certain way using policy-based routes — for example, to determine which 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.

SD-WAN Advantages

SD-WAN offers geographically distributed organizations and companies with multiple branches many benefits, including the following:

- » **Simplicity:** SD-WAN enables centralized management and simplified configuration rules. In addition, by 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) 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 (as shown in Figure 3-1). This can lead to inefficient resource usage 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:** Here are some ways that SD-WAN can lead to greater efficiency:
 - According to industry research, companies can save up to 40 percent over five years by cutting down on hardware, software, and support acquisition.
 - Fewer personnel are needed to manage, troubleshoot, and provision WAN equipment.

- 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.

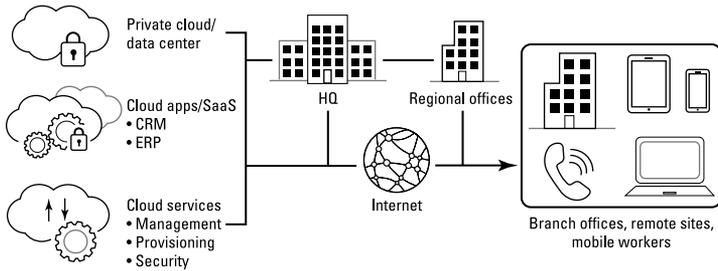


FIGURE 3-1: Efficient SD-WAN traffic routing.

Prioritizing Security



WARNING

When adopting SD-WAN, decision-makers often prioritize connectivity and cost benefits over security. This is a mistake that can put the network at risk.

Although SD-WAN offers many benefits, it can also bring challenges if it isn't architected correctly, 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. Plus, network performance can become 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 hubs 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. In addition, organizations can improve application performance when routing branch traffic to a distributed cloud edge versus a centralized hub. Leveraging the cloud for middle-mile connectivity can ensure greater end-user experience for the branch while also

diminishing the need to build a global backbone that's complex and time-consuming.



TIP

You may have already adopted SD-WAN in your network infrastructure (or you may be considering it) as a way to securely connect and control access to branch offices and remote employees. SASE creates a unified approach for SD-WAN and security services to connect to, providing a single point of view and simplified management solution to protect your network.

Understanding the Role of VPNs

For many years, VPNs have enabled secure connectivity to corporate networks and resources over the Internet. The two most common types of VPNs are remote access (for connecting remote users) and site-to-site (for connecting remote locations).

VPNs facilitate secure data transit over the Internet through a tunneling protocol, where data is 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.

VPNs are effective at enabling secure access to corporate data centers and other physical locations, but they aren't optimized for access to the cloud. As a result, there is no security or access control when users disconnect to reach cloud apps or services, as shown in Figure 3-2.

A SASE solution encompasses VPN services and enhances those capabilities. Operating in a cloud-based infrastructure, it securely routes traffic to physical locations, as well as to public cloud services such as SaaS, platform as a service (PaaS), infrastructure as a service (IaaS), and private cloud apps and services.

In an IPSec VPN, 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 (WAP), SD-WAN edge device, or firewall, all without the need to back-haul traffic to a physical location for security scrubbing (see Figure 3-3). Remote users can employ an always-on IPSec or SSL VPN connection between their endpoint or mobile devices and their applications, with the SASE solution ensuring consistent traffic encryption and threat prevention.

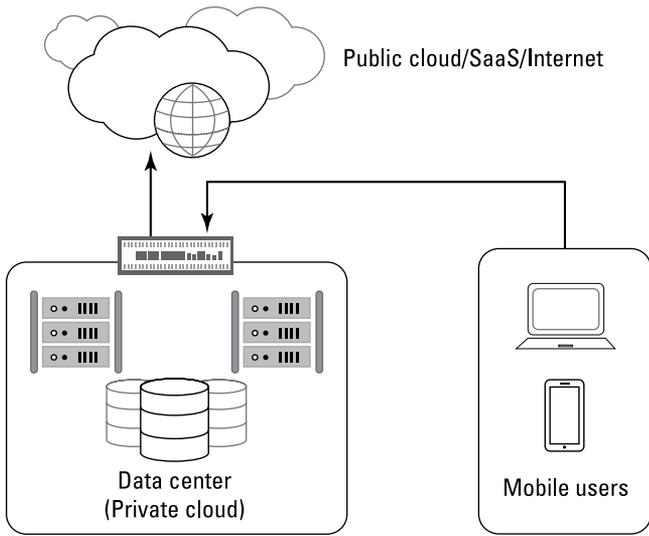


FIGURE 3-2: Remote-access VPN is not designed to support cloud applications.

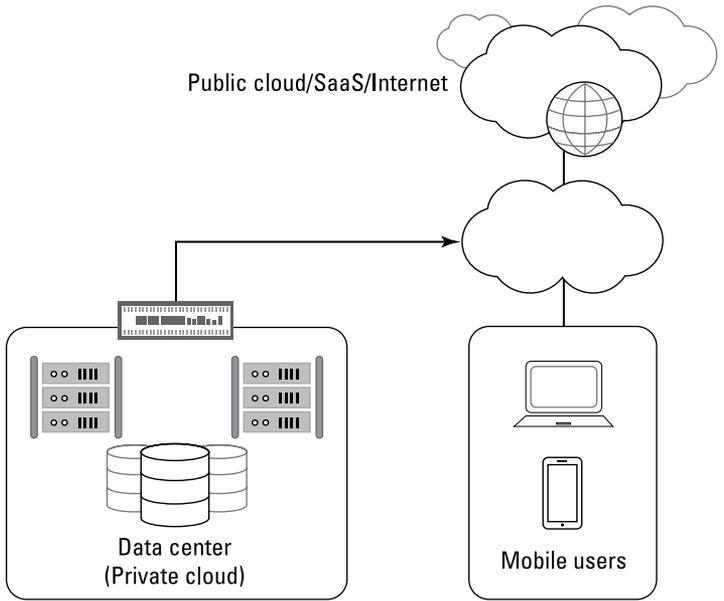


FIGURE 3-3: SASE uses cloud infrastructure to connect users to both cloud apps and the data center.



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 traffic to physical corporate locations. This dramatically simplifies the management and policy control needed to enforce least-privilege access rules.

Ensuring Quality of Service

As organizations transition from MPLS to SD-WAN using DIA links, they often find that the service quality varies. QoS controls establish bandwidth allocations assigned to particular apps and services and prioritize them when there is a contention for bandwidth. Businesses rely on QoS to ensure that their critical apps and services (for example, medical equipment or credit card processing services) perform adequately. QoS also helps businesses avoid and overcome bandwidth congestions caused by nonbusiness-related traffic, like streaming video, which can severely impact business operations and sales (see Figure 3-4).

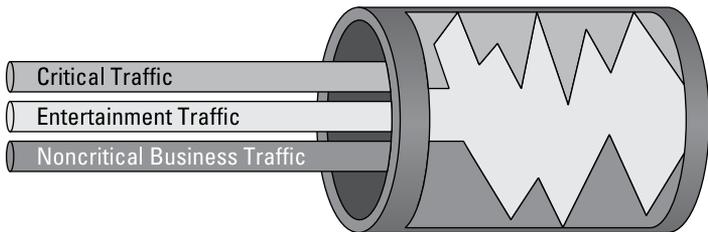


FIGURE 3-4: Bandwidth without QoS control.



WARNING

QoS is an important step when an organization begins migrating from MPLS to SD-WAN. As businesses start augmenting or replacing their MPLS links with broadband Internet, they realize it's a “best effort” that doesn't provide a service-level agreement (SLA) guaranteeing performance for application traffic. If you have QoS configured for your network, your broadband Internet service provider (ISP) will ignore QoS tagging on its routers. As a result, organizations typically procure enough bandwidth to meet their expected peak utilization, which can add to costs significantly.

Administrators can use QoS to designate which apps and services should take precedence over others, as shown in Figure 3-5. A SASE solution incorporates QoS services in the cloud, enabling IT administrators to easily mark sensitive applications, such as Voice over Internet Protocol (VoIP), as high priority over general Internet and entertainment sites and apps.

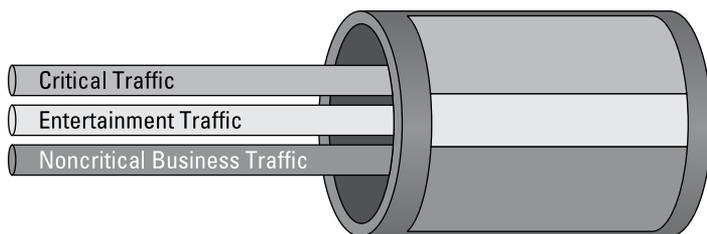


FIGURE 3-5: Bandwidth with QoS control.



REMEMBER

Managing the QoS traffic and allocation to guarantee performance SLAs doesn't need to be difficult. An effective SASE solution enables organizations to accurately identify applications and dynamically shape and prioritize traffic based on the business policies to better manage bandwidth and improve user experience.

Routing

Traditional architectures connected to wide area networks (WANs) uses routers to allow users to access applications, the cloud, and the Internet from local area networks (LANs). Routers connect networks, steer application traffic, and control user access on the network. As a result, these branch routers lack the ability to provide brownout controls and application remediation when WAN connections degrade or bandwidth saturation occurs. Businesses leveraging unified communications as a service (UCaaS) and SaaS applications with higher-bandwidth demands suffer from poor user experience due to such limitations and are forced to add more bandwidth to accommodate. Doing so can become costly.

Highly trained IT experts are required to configure branch routers. Using a command-line interface (CLI), they construct and configure core functionalities such as traffic forwarding, QoS, and access controls. When an organization has just a few routers, this may be workable, but it becomes a problem when an organization

has hundreds or even thousands of routers, dispersed across many branches worldwide. When organizations have service updates or need to upgrade their network, the idea of reconfiguring many routers manually just isn't practical.

Effective SASE solutions utilize intelligence and automation to avoid the inefficiencies of traditional routers. They simplify routing based on application intelligence, performance SLAs, and bandwidth requirements to accurately steer traffic. This results in effective bandwidth utilization on any WAN links (MPLS, broadband, direct Internet, and 4G/5G), which can decrease costs and improve application performance. Additionally, automating complex routing configurations and remediation tasks enables large-scale deployments while reducing troubleshooting and resolution efforts.

Accelerating SaaS

It's a no-brainer that organizations are moving to the cloud by adopting SaaS applications to simplify their branch infrastructures. As more and more branches require access to applications, it has become pivotal for businesses to address growing bandwidth demands and ensure exceptional user experience for employees at the branch. As a result, high jitter, latency, and packet loss can be detrimental to the application performance at the branch.



WARNING

Legacy networks, with their data center backhauling and packet-based routing, fail to intelligently steer traffic on the best-performing and highest available WAN links, while also adding significant latency to SaaS application access. Legacy approaches to SD-WAN have relied on taking the traditional model of packet routing and forcing it to fit the cloud-ready enterprise. Some solutions can appear to simplify the creation of VPNs over broadband connections but have fallen short in delivering on the transformative promise of SD-WAN.

Legacy SD-WAN solutions are built using Layer 3 packet-based policies, with limited app-based networking policies and app visibility. This makes it difficult for network teams to deliver application SLAs. As a result, businesses now need deep application visibility, with Layer 7 intelligence for network policy creation and traffic engineering. Only then can network teams

provide exceptional user experiences by delivering SLAs for all apps, including cloud, SaaS, and UCaaS.



REMEMBER

To combat these performance issues, a SASE solution can provide accurate application identification combined with advanced performance SLAs like mean opinion score (MOS), server response time, and transaction failures to steer SaaS traffic. In addition, they can automate application remediation to ensure consistent performance during network degradation. This results in critical improvements to user experience without added bandwidth requirements.

IN THIS CHAPTER

- » Implementing ZTNA
- » Ensuring Internet security with a cloud SWG
- » Identifying and securing access to SaaS apps
- » Deploying a next-generation FWaaS
- » Preventing sensitive data loss and ensuring regulatory compliance
- » Securing DNS resolution
- » Leveraging threat prevention tools

Chapter 4

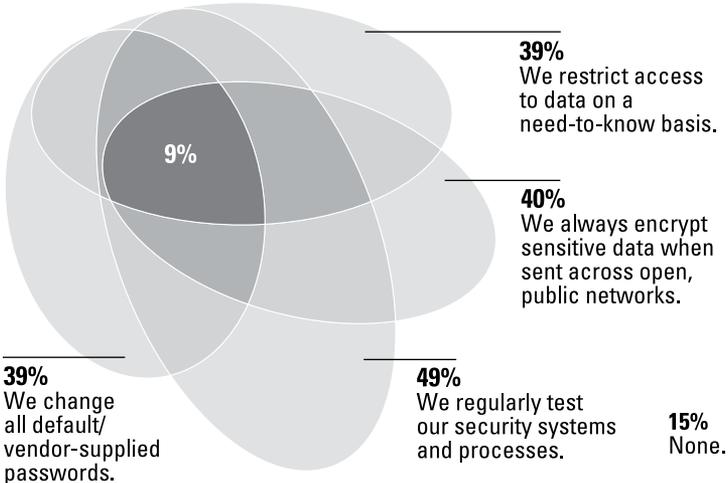
SASE Security Capabilities

In this chapter, you find out about the core security capabilities in a secure access service edge (SASE) solution. Security is a fundamental necessity for organizations as users connect to access data and applications from various locations, like the data center, headquarters, and branch offices. Here, you discover the critical tools needed to provide security across all different aspects of the network.

Modernizing the Access Infrastructure with ZTNA

Many companies still lack the necessary security protections and policies needed to adequately protect their users and data, as shown in Figure 4-1. This is especially true with the explosion

of the hybrid workforces demanding consistent and ubiquitous access to apps and data from everywhere. Legacy remote access solutions were never designed to facilitate or accommodate the massive shift away from the office and rapid rise of hybrid work. Having to work with systems ill-equipped for the current reality stresses IT security teams who are struggling to adopt their policies and techniques to this new corporate normal. As a result, more and more organizations are looking to modernize their access infrastructure around Zero Trust network access (ZTNA).



Source: Verizon Mobile Security Index 2021 report

FIGURE 4-1: Which of the following describe your organization's security policies?

ZTNA is a key part of the Zero Trust philosophy of “never trust, always verify,” developed by Forrester Research. Under ZTNA, users who want to connect to the cloud and access applications there must authenticate and have their traffic inspected up through Layer 7 via a gateway. This enables an IT admin to identify users and create policies to restrict access, minimize data loss, and quickly mitigate any issues or threats that may arise.

ZTNA solutions are based on a micro-perimeter architecture, ensuring least privileged access to authorized applications and data. ZTNA does not incorporate content inspection, and that creates a discrepancy in the types of protection available for each application. To deliver consistent protection, the organization must deploy additional controls on top of the ZTNA model to inspect all traffic across all applications. That ultimately leads to

increased deployment and management complexity and, thus, a less secure architecture.



REMEMBER

Layer 7 inspection and control, as well as advanced threat protection (ATP) security, are imperative to Zero Trust.

SASE builds upon the key principles of ZTNA and extends them across all the other services within a SASE solution. Identifying users, devices, and applications, no matter where they're connecting from, simplifies policy creation and management. 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 also support ZTNA capabilities for protecting applications, as well as incorporate advanced security services for the consistent enforcement of data loss prevention (DLP) and threat prevention policies. This is necessary because although access controls are useful for establishing who the person is, other security controls are also needed 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 to prevent 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. When they're 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.

Protecting Web Traffic with a Cloud SWG

Secure web gateways (SWGs) provide one solution to the problem of securing web traffic from a user endpoints (discussed in Chapters 2 and 3).

Instead of fully inspecting all network traffic, a web gateway examines traffic from a web browser and blocks websites and known malware. Organizations looking for a better solution (compared to having 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*, Google detected more than 2 million unsafe websites in January 2021.

SWG can be used to block inappropriate content (such as 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 granting Internet access.



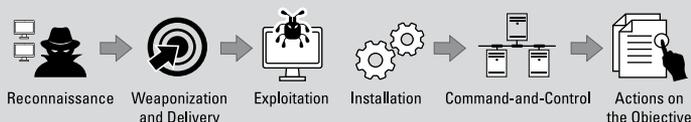
WARNING

Web gateways are not a substitute for firewalls. Partially inspecting traffic with an SWG 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, 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.

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

(continued)

(continued)

- 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 lifeblood 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. The attacker can also use this communication 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.



REMEMBER

A SASE solution should include SWG, enabling organizations to control web access and enforce security policies that protect users from hostile websites. But remember that a SWG is just one service of the overall SASE solution. Other security services like ZTNA, cloud access security broker (CASB), firewall as a service (FWaaS), Domain Name System (DNS) security, DLP, and ATP are also necessary to ensure all application traffic is secured.

Securing Access to SaaS Apps

SaaS applications (like Box, Microsoft 365, Microsoft Teams, Salesforce, and Slack) offer companies, employees, and customers many operational benefits. However, for each positive, there is also a negative when it comes to information security, as shown in Table 4-1.

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

TABLE 4-1 The Pros and Cons of SaaS Adoption

Pros	Cons
SaaS apps can be deployed quickly. As a software solution, SaaS app delivery and configuration is quick and painless. Apps are conveniently accessible directly to all users via the cloud from anywhere.	Anyone, including malicious actors, can use any cloud service. Cloud services are typically delivered without IT and security oversight. Users can access the application directly from both safe and unsafe connections anywhere and on any device — secure or not.
The cloud can store a lot of data at a low total cost of ownership (TCO) to the organization. Both users within the organization and external third parties can easily share cloud data.	SaaS application data 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.
SaaS apps are simple to maintain. Instead of having your IT department manually upgrade an app, that responsibility falls to the SaaS vendors, saving IT resources.	App maintenance isn't always for the purpose of increasing uptime. SaaS vendors do an amazing job of 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, at least in theory. The downside is that unsanctioned apps can grow virally in your organization.

As a result, some undesirable security tradeoffs have emerged:

- » **Lack of visibility** (and therefore protection) into sensitive data uploaded, created, and shared in the cloud.
- » **Direct access to applications and data from any device**, including unmanaged devices and bring your own device (BYOD) personal devices, and from anywhere, including unsafe public Wi-Fi in coffee shops or at home.
- » **Overexposure of sensitive data** through critical collaboration applications like Microsoft Teams, Slack, and Zoom. Data becomes difficult to protect when it consists of short, unstructured messages.



REMEMBER

» **Shadow IT**, in which employees use and access unsanctioned apps to get their work done if they aren't provided with the sanctioned IT tools that they need. Shadow IT introduces security risk and potentially exposes sensitive corporate data.

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. Unsanctioned SaaS applications are literally exploding in number, and it's nearly impossible to keep up with this growth in terms of user activity visibility and control.

Many organizations depend on 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 unintentional loss and threat actors. CASBs are cloud-based security policy enforcement points that provide a gateway for your SaaS provider and your employees.

The problem with legacy CASBs

Legacy CASB approaches to securing SaaS applications only partially address the modern problems, leaving organizations exposed and vulnerable due to several critical limitations:

» **They can't see and protect today's most critical apps.**

Legacy CASBs are focused only on scanning web traffic, which misses over half of all traffic, including traffic from non-web applications. They rely only on static databases and support requests for app discovery. This approach hinders their ability to identify or contain new SaaS applications before they're a risk. They lack the application programming interfaces (APIs) to secure the modern apps that distributed workforces heavily utilize.

» **They provide inaccurate data protection.** These solutions typically protect only data that goes through a proxy. They use separate tools and policies for SaaS and other control points, and often deliver inaccurate pattern-based detection that requires a ton of manual tuning. They weren't designed to detect sensitive natural language conversations embedded in messages on collaboration apps.

» **They traditionally put compliance before security.**

Security, unfortunately, has always been a checkbox in legacy CASB, with the majority of vendors using third-party or ineffective sandboxing as the only threat detection method. This approach provides very limited efficacy. As mentioned earlier, it gives visibility only into HTTP/S, leaving customers to find when something is missed via a high-priority threat or a breach.

In addition, legacy CASB products were designed with complex architectural constraints that are not effective anymore. They use a stand-alone proxy designed to perform a limited amount of inline inspection capabilities, in addition to API-based controls for introspective SaaS application security. Threat protection and DLP capabilities are siloed and not tightly integrated with the overall existing network implementation, requiring additional labor and expense to manage disjointed policies and multiple consoles and perform incident triage.

Next-generation CASBs



TIP

A next-generation CASB approach is natively 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 include the following:

» SaaS visibility

- Continuous identification of new apps
- Discovery and control of shadow IT usage
- Granular app risk assessment
- Configuration assessment

» Control and compliance

- Accurate data discovery and classification
- Advanced data discovery using techniques such as natural language processing, machine learning, and optical character recognition
- Compliance reporting and remediation
- Access control for managed and unmanaged devices

» SaaS protection

- ATP for known, unknown, and zero-day threats
- Data protection
- Workflow integration for incident auto-remediation

A SASE solution should natively incorporate next-generation CASB controls both in-line and in APIs to address the most stringent modern problem for SaaS application security and data protection.

A comprehensive SASE solution includes CASB and Enterprise DLP in a single integrated service, reducing network and security complexity while increasing organizational agility, assisting you through your cloud and network transformation while helping you safely adopt SaaS applications.

Deploying 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, 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 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 network 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 firewalls 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, and 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.



REMEMBER

A SASE solution should harness FWaaS capabilities to provide the protection of a next-generation firewall in the cloud. It's important to ensure your SASE solution doesn't just 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 DNS security.

Implementing Data Loss Prevention

Companies are processing massive amounts of data in more places than ever — in offices, in multiple SaaS applications, and in cloud storage environments. In addition, with cloud and mobile

computing technologies, employees can directly access applications and data anytime, anywhere, and from any device.

This ability to access data from anywhere introduces some data protection challenges, including the following:

- » The number of data breaches by insider threats continues to increase.
- » Most companies don't have much visibility into where their sensitive and regulated data is, or how and where their employees access it, use it, or share it with others. In cloud environments, that's even more of an issue.
- » SaaS and public cloud providers may offer some data protection capabilities, but they're often inadequate. This leads to ineffective and inconsistent security.

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 its sensitive data consistently across different repositories and communication vectors, such as Box, Microsoft 365, Slack, corporate devices, network traffic, and so on.
- » Monitor the 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.

- » 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. It consistently applies policies 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. This solution relies on a scalable and simple architecture and enables effective machine learning by leveraging access to all the organization's traffic and data.



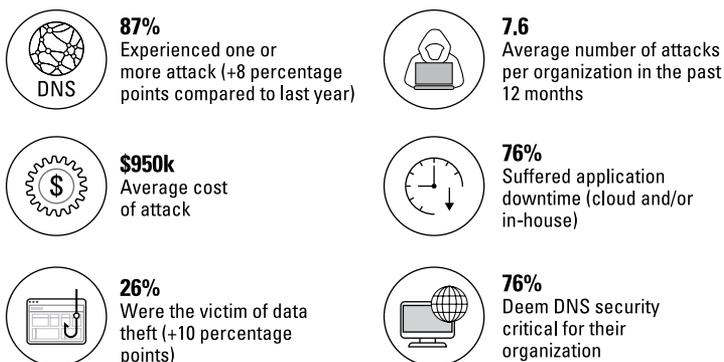
REMEMBER

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

Securing DNS

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 4-2).



Source: IDC 2021 Global DNS Threat Report

FIGURE 4-2: 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 85 percent of malware uses DNS to initiate C2 communications. 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

The SolarWinds supply chain attack became one of the most significant cybersecurity events at the end of 2020, impacting both commercial and government organizations worldwide. The DNS attack targeted SolarWinds Orion software with dormant malicious code that SolarWinds unknowingly sent out with a software update. The code resulted in hackers gaining access to IT systems, where more malware was then installed.

The attack used the SUNBURST Trojan, which uses DNS tunneling to receive commands from the adversary and exfiltrate data. The malware periodically contacted its C2 domain to report statuses and receive commands. When the C2 domain woke up from the incubation period, the majority of burst DNS requests were for new subdomains. The Trojan dynamically constructed these hostnames with *domain generation algorithms* (DGAs) to exfiltrate data. The DGA strings were encoded victims' identities, containing the infected organizations' domain names and security product statuses. When the attacker's DNS resolver received requests for these hostnames, it returned CNAME responses pointing to different C2 servers based on the exfiltrated information.

The SolarWinds supply chain attack leveraged DGA subdomains to exfiltrate data and provided a proxy layer for the attacking infrastructure.

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 DGA-based malware with instant enforcement.** Malware's use of 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 combats 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.



REMEMBER

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.

Protecting Networks from Threats

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. A variety of threat prevention tools exist, from anti-malware and intrusion prevention to Secure Sockets Layer (SSL) decryption and file blocking, providing organizations ways to block threats. However, these point products require separate solutions, making management and integration difficult.

A SASE solution integrates all these point products and services 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's being provided by the vendor. The vendor should gather and share 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
- »» ATP 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

IN THIS CHAPTER

- » Understanding the user experience challenges organizations face
- » Identifying the key requirements for user experience monitoring
- » Realizing the value of SASE native digital experience monitoring

Chapter 5

Digital Experience Monitoring

In this chapter, you find out about the user experience monitoring capabilities of a secure access service edge (SASE) solution. Security and connectivity are important, but having exceptional user experience is key because it impacts productivity and operations.

User Experience Challenges

The concept of the “corporate network” has greatly expanded, providing more work for IT teams and increased opportunity for employees to become frustrated when they can’t access the tools they need to do their jobs. Employees need consistency in both security and user experience as they move among branch, home, and other locations. IT teams need complete visibility into their workplace’s end-user experience in order to support employees when performance problems arise.

User-experience challenges many businesses face include the following:

- » **Security and user experience are seen as a tradeoff.** Legacy networking and security architectures rely on backhauling all traffic to corporate data centers. This forces network administrators to choose between security and performance for their users.
- » **There's a gap between what users experience and what IT sees.** With corporate resources distributed across public clouds, software as a service (SaaS), and corporate data centers, IT teams struggle to identify and diagnose application performance degradation. The hybrid workplace presents additional challenges as the plethora of home routers, Wi-Fi networks, and Internet service providers (ISPs) introduce additional points of impact where IT has no visibility and control.
- » **Existing monitoring approaches aren't SASE native and leave you in the dark.** Legacy monitoring solutions weren't designed for a hybrid workplace. They provide only a fraction of the visibility needed into true end-user experience and often require additional software or hardware.

Together, these factors have contributed to costly troubleshooting processes, loss of productivity, and poor user experiences in today's hybrid workplace. Digital experience monitoring solutions used for network, endpoint, and application monitoring provide siloed visibility into their respective domains but lack the context of the overall SASE environment, making it difficult to troubleshoot effectively.

According to a recent ESG report, organizations are beginning, in process, or at mature digital transformation initiatives (88 percent), with creating a differentiated experience (40 percent) as a top goal of digital transformation (see Figure 5-1).

Managing the digital experience

Digital experience monitoring (DEM) has emerged as a new Gartner IT category to address user experience, human or machine, across every dependency, whether network or service, inside or outside your company.

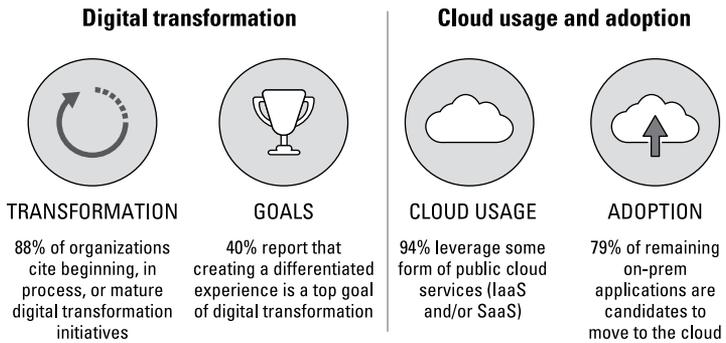


FIGURE 5-1: Palo Alto Networks, please provide a caption here.

DEM technologies monitor the availability, performance, and quality of experience an end user or digital agent receives as they interact with an application and the supporting infrastructure. Users can be external consumers of a service (such as patrons of a retail website), internal employees accessing corporate tools (such as a benefits management system), or a combination of both. DEM technologies seek to observe and model the behavior of users as a continuous flow of interactions in the form of user journeys.

But digital experience increasingly depends on a host of external services (like cloud, SaaS, and the Internet) that you don't own or directly control. You may not own all the underlying infrastructure, but you still own user experience. A comprehensive DEM solution includes the following monitoring approaches:

- » **Endpoint monitoring:** Includes the collection of information about things like CPU utilization, memory usage, and Wi-Fi signal strength to determine the negative impact those things may be having on a user's digital experience.
- » **Real-user monitoring (RUM):** Tracks performance based on data coming from actual users. It's a reliable technique to monitor how the application is being used and how real-world parameters such as network latency, device variation, and so on affect the end-user experience.
- » **Synthetic monitoring:** Includes running regular tests from a source to a destination, or from a user to an application, that allow IT to monitor network and application service performance even when they aren't being used. Synthetic

monitoring enables IT to understand the performance of assets they don't own and can't instrument, like SaaS applications and user ISPs. With synthetic monitoring, IT can easily diagnose a problem, because they have a baseline of app performance that enables them to pinpoint when and where the performance bottleneck is, like a network outage. Synthetic monitoring is one of the most critical technical components of DEM.

Automating DEM

Gartner predicts that by 2025, 70 percent of digital business initiatives will require infrastructure and operations (I&O) leaders to report on business metrics from digital experience, up from less than 15 percent today.

When *autonomous digital experience management* (ADEM) capabilities are integrated into a SASE solution, even the most novice of analysts can derive deep insights into application performance without advanced IT specialization and separate management on endpoints and branch devices.



REMEMBER

ADEM enhances your IT team's tasks with easy-to-use, single-pane visibility that leverages endpoint, simulated, and real-time user traffic data to provide the most complete picture of user traffic flows possible.

Modern ADEM should provide business with the following:

- » **SASE-native DEM:** DEM capabilities are natively integrated into a SASE solution to optimize experiences for every user, working from anywhere, without the complexity of installing additional software or hardware.
- » **Segment-wise insights:** Operators can view every segment in the application delivery path for all users — in a branch office or at home — to help find root causes fast and expedite troubleshooting.
- » **Comprehensive visibility:** A unified view into the entire user experience journey requires performance insights from endpoint devices, networks and applications — all from one dashboard. Only with these insights will you have a true end-to-end view of everything that affects digital employee experience and productivity.

With ADEM, IT and security teams have the advantage of centrally implementing and monitoring remote access security policies and user experience for their hybrid workforces through a single pane of glass. And when capabilities are natively integrated with a SASE solution, no additional agents are required and no additional burden is placed on the user.



WARNING

Don't be fooled by complex monitoring solutions that require additional agents to be installed on end-user devices.

Identifying the key benefits of ADEM

The best business decisions are made when an organization views and understands its service from the end user's point of view. ADEM helps quickly pinpoint issues and deliver great digital experiences for every user, so IT can stop guessing what the problems are and start optimizing.



TIP

Stop guessing, and start optimizing. Use a unified platform, with a single pane of glass, to give teams the precise, automated insights and context they need to proactively deliver better user experiences and drive better business outcomes.

SASE native digital experience monitoring offers companies the following benefits:

- » **An integrated solution:** Organizations can have a fully integrated SASE solution with unified management of both mobile users and remote/branch networks without adding additional software or hardware.
- » **Reduced ticket volume:** Proactive, synthetic monitoring approaches, in addition to real user visibility, help address tickets before they're reported.
- » **Reduced time to isolate problems:** A single dashboard for network, application, and desktop support lets operators quickly see and resolve user problems with precision, regardless of their location.
- » **Increased efficiency:** With visibility and the right insights to quickly determine root causes, teams can stop finger-pointing and wasting hours on issue validation.

Seeing How SASE Native Monitoring Adds Value

Monitoring provides value across the organization, assisting both IT teams and end users — from outages outside an organization to common user device problems. With user experience monitoring, there is no longer any need for IT to rely on manual processes to identify and mitigate potential issues or use monitoring methods that only provide a fraction of the required visibility needed to troubleshoot effectively. Here are a few examples of how SASE native digital experience monitoring can add value.

Quickly identifying and resolving end-user device issues

With a single pane of glass into the SASE environment, IT can quickly identify and locate an issue, narrowing it down to an individual user device. For example, a user working at home using their home Wi-Fi may end up moving around in their house. In some areas of the house, the user may encounter low signal strength, which in turn negatively impacts their application experience. Another scenario might find a user performing an operating system upgrade, or an upgrade starting without the user's knowledge. In either case, this event may result in high CPU usage. Monitoring capabilities would be able to observe that a user's new operating system (OS) software install/update resulted in high CPU usage, which in turn would negatively impact performance on user applications.

Possessing this level of visibility ensures that operations teams can significantly reduce the amount of time needed to locate and fix an issue that's impacting user experience. IT can determine if Internet connections or endpoint devices are causing issues and proactively notify users of the potential problem.

Optimizing the hybrid workplace experience

Now that hybrid workplace is the norm, companies will need to continuously monitor every user's experience to make sure their user experience is always consistent as they shift between working from home, connecting over nontrusted networks, and

working from the office over a trusted corporate campus network. Monitoring capabilities offer deep insights and visibility into every part of the service delivery chain impacting user experience. That monitoring can include watching for device issues, home Wi-Fi, and network issues, Internet path issues, and issues with the applications itself. All that information can enable IT to quickly isolate problems and resolve issues for any user, working from anywhere.

Monitoring the branch experience

Operators need not only a unified view into user experience, but also visibility into remote offices located across the world. When monitoring capabilities are natively integrated into a software-defined wide area network (SD-WAN), organizations can monitor end-to-end user experiences for critical branch endpoints, including Internet of Things (IoT) devices. This enables an administrator to view an application experience score on a per-path basis, whether that path is active or backup. When they're able to run proactive synthetics on every path, an administrator can recognize the best path, per application, for all users in a branch office. Without native integration into an SD-WAN, IT wouldn't be able to attain such comprehensive visibility.

IN THIS CHAPTER

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

Chapter 6

Ten Benefits of SASE

Thinking that secure access service edge (SASE) may be the right choice for your business? Here are ten important business and technical benefits of deploying 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 users — no matter their location.

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

Greater Control of Users, Data, and Apps

Users are increasingly leveraging a variety of applications — including software as a service (SaaS) applications from multiple devices and locations — for both work-related and personal purposes. Many applications, such as instant messaging (IM), peer-to-peer (P2P) file sharing, and Voice over Internet Protocol (VoIP), can operate on nonstandard ports or hopping ports. Some of these applications may be sanctioned by the organization, others tolerated, and still others 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 to configure appropriate policies and rules. SASE provides complete visibility into application usage along with capabilities to understand and control their use, as shown in Figure 6-1.

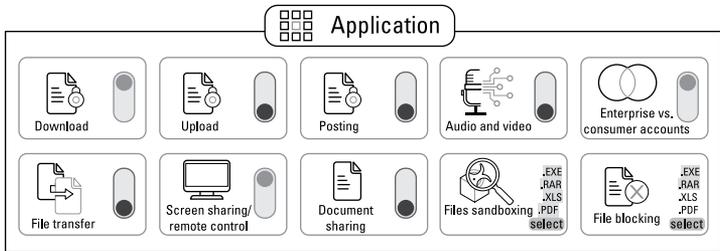


FIGURE 6-1: Control application usage in policy.

Better Monitoring and Reporting

SASE eliminates the need to monitor multiple consoles across different networking and security products and creates 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. The goal is to eliminate disjointed data protection policy and configurations for different SaaS apps, for on-premises repositories, and so on. Those things too often cause security blind spots, complex manageability, policy inconsistency, shadow IT, and shadow data. Instead, SASE enables a consistent data loss prevention (DLP) policy across every environment where data lives and flows, regardless of its location. You can rapidly and easily deploy new security services and applications with specific security policies, from the cloud to branch and retail locations, instead of having to individually manage them at each location.

Reduced 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 many 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 can 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. Chapters 3 and 4 explain 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 software-defined wide area network (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 configured for failover. This helps reduce congestion and latency associated with backhauling Internet traffic across MPLS connections or routing traffic across a connection that's experiencing high utilization or performance issues.

Enhanced User Experience

User experience is key for productivity and employee satisfaction. Digital experience monitoring (DEM) helps identify and remediate user-experience problems before they impact employees, IT, and the business. With SASE, DEM improves operations and optimizes experiences for every user, working from home or from branch offices, without the complexity of installing additional software and hardware.

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.

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.

autonomous digital experience management (ADEM): A digital experience management solution that offers SASE-native visibility into digital experience, including segment-wise insights across the entire service delivery path.

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.

command-line interface (CLI): A text-based user interface used to run programs, manage computer files, and interact with the computer.

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

digital experience monitoring (DEM): The process of optimizing the operational experience and behavior of a user by viewing the application and service portfolio.

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.

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

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

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.

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.

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 service provided maintains the underlying physical cloud infrastructure.

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

intellectual property (IP): Owned information, including 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 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.

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.

mean opinion score (MOS): A way to measure the quality of experience across video, audio, or other telecommunications.

multicloud: 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.

personally identifiable information (PII): Data (such as name, address, Social Security number, birth date, 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 (such as a financial institution, online auction site, retail store, 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 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 *and* User Datagram Protocol.

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, secure web gateway, cloud access security broker, firewall as a service, *and* Zero Trust network access.

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.

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 in 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, and .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): The unique identifier for any resource connected to the web. Commonly known as a *web address*.

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): A telephony protocol that is 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.

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.

Notes

Notes

Notes

Notes



Connect and Secure Without Compromise

Complete, best-in-class security with exceptional user experience

Palo Alto Networks Prisma® SASE converges best-in-class security with best-of-breed Next-Gen SD-WAN into a single cloud-delivered service. It consolidates multiple point products, including ZTNA, Cloud SWG, CASB, FWaaS, and SD-WAN, reducing network and security complexity while increasing organizational agility.

Learn how Prisma SASE can secure your hybrid workforce, whether users are remote, mobile or working from an office.

<https://www.paloaltonetworks.com/sase>



Build your digital transformation on a SASE framework

Enabling a secure hybrid workforce and supporting digital transformation have become top priorities for organizations everywhere. However, traditional network and security architectures weren't designed with the cloud or hybrid work in mind, creating problems with complexity, visibility, flexibility, end-user experience, and incomplete protection. A secure access service edge (SASE) provides connectivity, consistent security, and optimized user experience for your hybrid workforce, 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: 978-1-119-89742-2

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.