

LEARNING MADE EASY

Palo Alto Networks Special Edition

Zero Trust Network Access

for
dummies[®]
A Wiley Brand



See why
VPN is dead

Enable true least-privileged
access with ZTNA 2.0

Understand legacy
ZTNA limitations

Brought to you
by



Lawrence Miller

About Palo Alto Networks

Palo Alto Networks is the world's cybersecurity leader. We innovate to outpace cyberthreats, so organizations can embrace technology with confidence. We provide next-gen cybersecurity to thousands of customers globally, across all sectors. Our best-in-class cybersecurity platforms and services are backed by industry-leading threat intelligence and strengthened by state-of-the-art automation. Whether deploying our products to enable the Zero Trust enterprise, responding to a security incident, or partnering to deliver better security outcomes through a world-class partner ecosystem, we're committed to helping ensure each day is safer than the one before. It's what makes us the cybersecurity partner of choice.

At Palo Alto Networks, we're committed to bringing together the very best people in service of our mission, so we're also proud to be the cybersecurity workplace of choice, recognized among Newsweek's Most Loved Workplaces (2021), Comparably Best Companies for Diversity (2021), and HRC Best Places for LGBTQ Equality (2022). For more information, visit www.paloaltonetworks.com.

Zero Trust Network Access

for
dummies[®]
A Wiley Brand



Zero Trust Network Access

Palo Alto Networks Special Edition

by Lawrence Miller

**for
dummies**[®]
A Wiley Brand

Zero Trust Network Access For Dummies®, Palo Alto Networks 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.

ISBN 978-1-119-98379-8 (pbk); ISBN 978-1-119-98381-1 (ebk)

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.

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

Client Account Manager:
Cynthia Tweed

Production Editor:

Magesh Elangovan

Special Help: Don Meyer,
Shannon Bonfiglio

Table of Contents

INTRODUCTION	1
About This Book	2
Foolish Assumptions.....	2
Icons Used in This Book.....	3
Beyond the Book.....	3
CHAPTER 1: Recognizing the Security Implications of the New Normal.....	5
Looking at the Changing Landscape	5
Threats increasing in sophistication and frequency.....	6
Too many tools and too much complexity	6
Shortage of cybersecurity talent and skills.....	7
Understanding the Need for Change.....	7
Evolution of work from a place we go to an activity we perform	8
Users everywhere, apps everywhere	9
Direct-to-app connectivity exponentially increases your attack surface	10
VPNs are too coarse	11
What Is Zero Trust Network Access?.....	11
The basics of ZTNA	11
ZTNA 1.0.....	12
ZTNA 1.0 Has Major Limitations in Today's Environment	13
Violates principle of least privilege	13
Incorporates an allow-and-ignore model	14
Does not provide security inspection.....	15
Does not protect data	16
Does not secure all apps.....	16
CHAPTER 2: Introducing Zero Trust Network Access 2.0.....	19
Fully Realizing Least-Privilege Access.....	19
Enabling Continuous Trust Verification	20
Ensuring Continuous Security Inspection	21
Protecting All Data.....	22
Securing All Apps.....	23

CHAPTER 3:	Understanding Critical Capabilities for ZTNA 2.0 Success	25
	Delivering an Exceptional User Experience.....	25
	Providing a Unified Solution	26
CHAPTER 4:	How to Get Started with ZTNA 2.0	29
	VPN Replacement.....	29
	Securing Internet Access	34
	Advanced SaaS Security.....	37
CHAPTER 5:	Ten (or So) Questions to Ask Your ZTNA 2.0 Vendor	41
	Do You Provide Full Layer 7 Application Visibility?.....	41
	Do You Provide Continuous Trust Verification?	42
	Do You Consistently Secure All Apps in a Single Product?.....	43
	Do You Apply Complete Security Inspection?	43
	Do You Consistently Protect All Enterprise Data?	43
	Do You Provide Uptime and Performance SLAs for All Apps?	44
	Do You Have a Single Unified Product to Secure the Enterprise?.....	44
	GLOSSARY	45

Introduction

How and where we work has changed dramatically in a relatively short period of time. Digital transformation initiatives that were already underway before the COVID-19 pandemic, such as remote working and cloud computing, were suddenly and necessarily accelerated to address new realities in the modern world. We now live in a world where work is no longer a place we go. Instead, it's an activity we perform anywhere.

As a result of this new work paradigm, our attack surface has increased exponentially, with many architectures now supporting direct-to-app connections across the Internet instead of back-hauling traffic to data centers over private networks. Legacy remote access VPN connectivity no longer works in a world where users and apps are now outside corporate networks and data centers. Legacy remote-access VPNs provide too much access with little to no threat or vulnerability detection, leaving privileged resources vulnerable to user account compromise. With the dramatic increase in volume, scale, and sophistication of cyberattacks, today's cloud-enabled businesses are scrambling to plug their security 'gaps' and have started turning to zero trust network access (ZTNA) solutions to reduce their attack surface and protect their enterprises from ransomware and other exploits.

However, existing ZTNA (or 1.0) solutions are unable to meet the security needs of today's enterprises. They provide too much access with too little protection, deliver inconsistent and incomplete security across web- and non-web-based applications, and offer poor performance and user experiences. As a result, they are not able to handle the onslaught of new and increasingly sophisticated attacks across our exploding attack surfaces.

ZTNA 2.0 solutions have emerged as the best path forward, ushering in a new era of secure access in a world where work is an activity, not a place.

About This Book

Zero Trust Network Access For Dummies, Palo Alto Networks Special Edition, consists of five chapters that explore the following:

- » The changing security landscape, the basics of ZTNA, and the need to move beyond ZTNA 1.0 (Chapter 1)
- » How ZTNA 2.0 addresses the limitations of current ZTNA solutions (Chapter 2)
- » The critical success factors to look for in a ZTNA 2.0 solution (Chapter 3)
- » Key ZTNA 2.0 use cases and customer success stories (Chapter 4)
- » Important questions to ask your ZTNA vendor (Chapter 5)

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

There's also a helpful glossary in case you get stumped by any terms or acronyms used in this book.

Foolish Assumptions

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

Mainly, I assume that you're a technology decision-maker or practitioner and you're looking for an innovative solution to deliver secure access for your hybrid workforce. Whether you're a chief information security officer (CISO), an IT manager, or a network or security engineer, this book illustrates how ZTNA 2.0 can help you address the challenges of a greatly expanded attack surface and an increasingly hostile threat landscape.

Icons Used in This Book

Throughout this book, I 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.



TECHNICAL
STUFF

This icon explains the jargon beneath the jargon and is the stuff legends — well, legendary nerds — are made of.



TIP

Tips are appreciated, but 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.

Beyond the Book

There's only so much I can cover in this short book, so if you find yourself at the end of this book wondering, "Where can I learn more?," go to www.paloaltonetworks.com/sase/ztna.

IN THIS CHAPTER

- » Surveying the evolving security landscape
- » Recognizing the changing nature of work
- » Understanding the basics of Zero Trust
- » Looking at the limitations of ZTNA 1.0

Chapter 1

Recognizing the Security Implications of the New Normal

This chapter explores modern security challenges including increasing threats, complexity in the security ecosystem, and the cybersecurity talent shortage. It also explains the basics of Zero Trust Network Access (ZTNA) and why organizations today need to adapt their remote access strategies to align with new work models and evolve beyond traditional access control solutions.

Looking at the Changing Landscape

The modern security landscape continues to evolve as threats have increased in both sophistication and frequency. In response to these threats, organizations have deployed an ever-growing and dizzying array of point security solutions and tools. However, managing and operating these siloed tools often requires specialized skills and resources that most enterprise security teams simply don't have.

Threats increasing in sophistication and frequency

Data breaches and ransomware attacks have become so frequent today that they practically warrant their own news segment alongside weather, sports, and traffic. The fact that these security events are commonplace, however, does not make them less dangerous. Organizations that become complacent in their security posture risk extensive damage when an attack occurs.



WARNING

According to the Ponemon Institute, the average cost of a data breach increased by 10 percent to \$4.24 million from 2020 to 2021. This was the largest single-year cost increase in the last seven years.

Unfortunately, enterprise security teams are fighting an uphill battle in the face of increasingly advanced tactics, techniques, and procedures (TTPs) used by threat actors.



REMEMBER

An organization's ability to stay on top of the modern threat landscape requires effective tools and a team of capable security analysts. Unfortunately, having the proper balance of technology and skilled experts tends to be the exception for most organizations, rather than the rule.

Too many tools and too much complexity

Enterprise security teams have been deploying “one-off” point security solutions to address specific security challenges and limited use cases for many years. This has often been mistakenly rationalized as “defense in depth.” The unfortunate result is that the security ecosystem is littered with too many tools that create a complex, costly, and ineffective operating environment. According to a 2020 IBM study, the average enterprise uses 45 security tools, and 30 percent of organizations use more than 50 tools. According to the *Panaseer 2022 Security Leaders Peer Report* cited in *InfoSecurity Magazine*, “the shift to cloud and remote working has driven a 19 percent increase over the past two years in the number of security tools organizations must manage — from 64 to 76.”

These security tools typically generate thousands of alerts every day — far exceeding the volume that enterprise security teams

are staffed to effectively handle. The alerts come from many disconnected tools, leaving security analysts to piece together the puzzle (see Figure 1-1).




FIGURE 1-1: Too many security tools results in complexity and alert fatigue.

Shortage of cybersecurity talent and skills

Beyond the increasing frequency and sophistication of threats and the growing complexity and number of security tools in the enterprise, the challenges of the modern security landscape are further exacerbated by the global shortage of cybersecurity talent and skills. The Information Systems Audit and Control Association (ISACA) estimates that almost two-thirds of enterprise security teams are understaffed, and more than half have open positions. The International Information System Security Certification Consortium (ISC)² estimated the global shortage of cybersecurity professionals to be 2.72 million in 2021.

Understanding the Need for Change

In addition to a rapidly evolving threat and security landscape, changes in the nature of work and how and where applications and data are accessed are driving the need for fundamental changes to the notion of trust and how we grant users and devices access to our applications and data.

Evolution of work from a place we go to an activity we perform

The nature of work has changed — from a place workers go to an activity that they perform. We no longer “go to work”; now we just “work.” For many businesses, the location of their workers and where their individual work duties are performed has become largely irrelevant. We’re now able to perform our activities whenever and wherever we need to. This change is driven by two important trends:

» **Applications are everywhere.** Most enterprises have moved to a model where you no longer consume applications running in a corporate data center. The application delivery model — including software as a service (SaaS), web, and cloud — is now hybrid. The overwhelming majority of enterprises today utilize some combination of private cloud, public cloud, Internet, and SaaS.

According to the Flexera *2021 State of the Cloud Report*, 80 percent of organizations have a hybrid cloud strategy. Statista reports that the average organization uses 110 SaaS applications.



TIP

» **Users are everywhere.** Many organizations today have adopted a hybrid work model, supporting *partially remote* (working out of a home office two to three days a week), completely remote, or anything in between. This trend was greatly accelerated by the global pandemic and, as companies have realized productivity and employee morale benefits, has become the new normal for work.

According to *The State of Hybrid Workforce Security 2021* report from Palo Alto Networks, 76 percent of employees want to be hybrid, even after the pandemic.



TIP

However, this change has some important ramifications for IT and security.

Previously, organizations connected their remote workforces to data centers and protected access to the applications in the data center and any web and SaaS apps. This was achieved by deploying various point security tools at the data center perimeter: firewalls, proxies, intrusion prevention systems (IPSs), cloud access

security brokers (CASBs), anti-malware protection, Domain Name System (DNS) security, and so on.

In this model, organizations built their wide-area networks (WANs) with Multiprotocol Label Switching (MPLS) and other dedicated links connecting the branch offices to the data center. All Internet-bound traffic was routed through the data center, which meant that this massive security stack could be centralized and all traffic sent through it (see Figure 1-2).

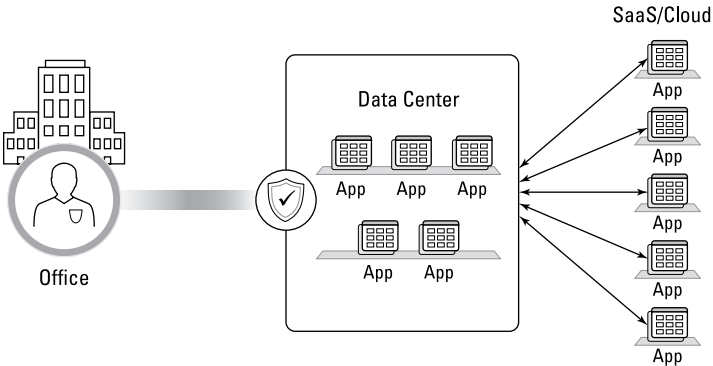


FIGURE 1-2: Security was relatively simple when work was a place you went to every day.

What we’re now seeing is a completely different model.

Users everywhere, apps everywhere, data everywhere

Organizations have shifted their WAN architecture from connecting remote workforces to data centers to now going directly to the Internet. As a result, they must now focus on delivering secure and reliable access to users working from anywhere — from corporate and branch offices, from home offices, and on mobile devices — and connecting to applications and data that are located everywhere — in data centers, private clouds, public clouds, and SaaS apps.

Now users connect directly to all the applications needed to work (see Figure 1-3). The location of the application matters less; what matters now is delivering a consistent, optimized, and secure experience to access all those applications.




FIGURE 1-3: Users are now connecting directly to their apps.

Direct-to-app connectivity exponentially increases your attack surface

This direct-to-app connectivity is a dramatic shift from the traditional model, and it exponentially increases the enterprise attack surface. The more the attack surface expands, the more you need both security and access controls to protect enterprise applications and data (see Figure 1-4).




FIGURE 1-4: The attack surface has exploded.

VPNs are too coarse

Virtual private networks (VPNs) were designed to grant access to a local-area network (LAN) or subnet within the LAN, offering a private, encrypted tunnel for remote employees to connect to the corporate network. Although this may seem like a practical solution, VPNs unfortunately lack the flexibility and granularity to control and see exactly what users can do and which apps they can access. After a user is granted access, they can access anything on the network or subnet, leading to security gaps and policy enforcement problems.

ZTNA, on the other hand, provides secure remote access to applications based on granular access control policies. It offers access only to authorized apps for users instead of the “after you’re verified, you can access anything” approach that VPNs take. Thus, ZTNA provides a least-privilege approach to dramatically reduce the attack surface and improve the overall security posture.

What Is Zero Trust Network Access?

ZTNA is a category of products that provides secure remote access to applications and services based on defined access control policies. ZTNA solutions default to deny, providing only the access to an app or service the user has been explicitly granted. It’s important to understand the security gaps and benefits ZTNA solutions can provide organizations as more remote users join the network.



REMEMBER

ZTNA is a key part of the Zero Trust philosophy of “never trust, always verify,” developed by Forrester to identify the need to protect data. ZTNA requires users who want access to applications to authentication through a gateway broker before gaining access to the applications they need. This requirement provides the ability to identify users and create policies to restrict access, minimize data loss, and quickly mitigate any issues or threats that may arise.

The basics of ZTNA

With ZTNA, access is established after the user has been authenticated through the access broker. The ZTNA service then provisions access to the application on the user’s behalf through

a secure, encrypted tunnel. This provides an added layer of protection for corporate applications and services by shielding otherwise publicly visible Internet Protocol (IP) addresses.

Like software-defined perimeters (SDPs), ZTNA leverages the concept of a “dark cloud,” preventing users from seeing any applications and services that they don’t have permission to access. This protects against lateral movement, where a compromised endpoint or credentials would otherwise permit scanning and pivoting to other services by an attacker.

ZTNA 1.0

Initial ZTNA, or ZTNA 1.0, solutions were introduced at a time when the threat landscape, corporate networks, and how and where people worked were vastly different than they are today. As a result, ZTNA 1.0 solutions no longer align with the new world of work, and malicious actors are finding new ways to exploit the limitations of these ZTNA 1.0 approaches.

ZTNA 1.0 was designed to protect organizations by limiting their exposure and reducing their attack surface. It leverages an access broker to facilitate connectivity to an application. When a user requests access to an application, the access broker determines whether the user should have permission to access an application. After the permission is verified, the access broker grants access, and the connection is established (see Figure 1-5).

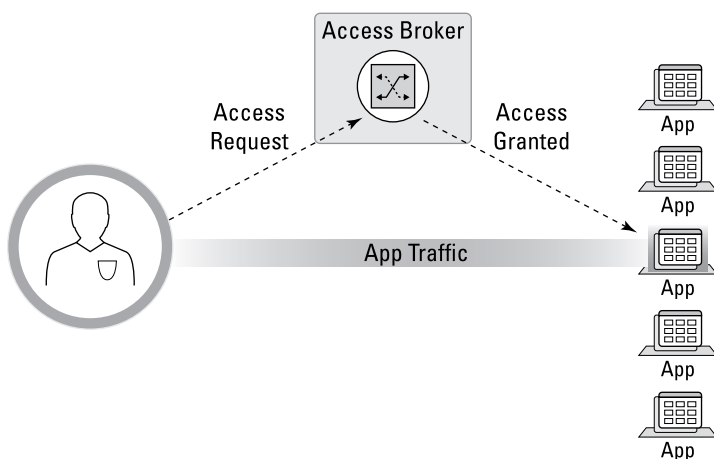


FIGURE 1-5: The industry tried to solve secure access with ZTNA 1.0.

And that's it. The broker is no longer in the picture, and the user is now given complete access to the application without any additional monitoring from the security system.

This “allow and ignore” approach is the architectural model of ZTNA 1.0. This model isn't just problematic in the context of today's threat landscape — it's dangerous.

ZTNA 1.0 Has Major Limitations in Today's Environment

Many ZTNA 1.0 solutions are based on SDP architectures, which do not provide content inspection, thus creating a discrepancy in the types of protection available for each application. In terms of consistent protection, it is up to the organization to incorporate additional controls on top of the ZTNA model and establish inspection for all traffic across all applications. Beyond these challenges, there are five major issues associated with ZTNA 1.0 solutions that limit their effectiveness in today's rapidly evolving security landscape and work environment.

Violates principle of least privilege

The principle of least privilege requires that a user is granted only the minimum level of access to an application or resource necessary to perform an authorized task, and nothing else. A Zero Trust strategy implies that nothing attempting to connect to a network application or resource — including users, applications, and devices — is inherently trusted.

Existing ZTNA 1.0 solutions manage application access at Layers 3 (Network) and 4 (Transport) of the Open Systems Interconnection (OSI) model using only IP address and Transmission Control Protocol (TCP) and User Datagram Protocol (UDP) port constructs.

A network is not the same as an application, but ZTNA 1.0 solutions rely on network-level access controls to provide users with application-level access. Unfortunately, relying on policies defined at Layers 3 and 4 creates a number of problems. For example, if an app uses dynamic ports or IP addresses, you must grant access to broad ranges of IP addresses and ports, exposing

more surface area than necessary. Access can't be restricted at the sub-app level or app-function level either; access can only be granted to entire apps. The inevitable result is that users end up with far more access than desired or intended (see Figure 1-6).

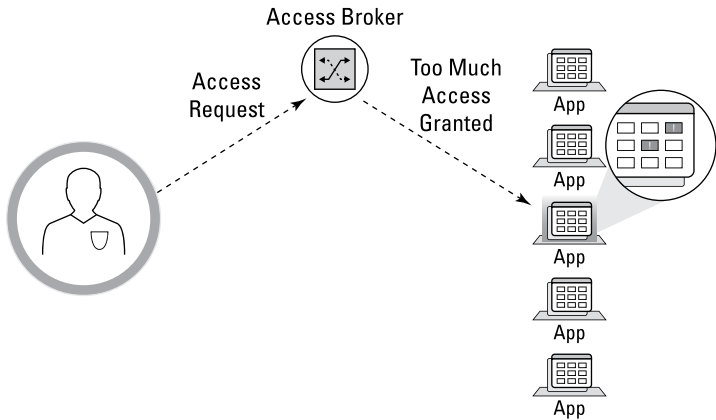


FIGURE 1-6: ZTNA 1.0 violates the principle of least privilege.



WARNING

Any malware that listens on the same IP addresses and port numbers as allowed applications can freely communicate with command-and-control (C2) infrastructure and spread laterally.

Incorporates an allow-and-ignore model

Another limitation of ZTNA 1.0 solutions is that they rely on a risky allow-and-ignore model (see Figure 1-7). When the access broker establishes the connection between the user and the application, user and device traffic is trusted and no further verification is performed for the duration of the session.

Assuming that trust only needs to be verified once and never checked again is a recipe for disaster. A lot can happen after trust is initially established. User and application behavior can change, and applications can be compromised.



WARNING

Many modern threats piggyback on allowed activity to avoid triggering alarms.

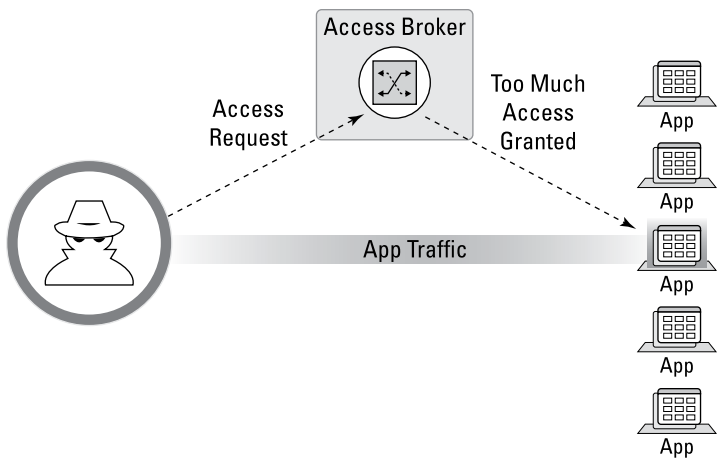


FIGURE 1-7: ZTNA 1.0 allows and ignores.

Does not provide security inspection

ZTNA 1.0 solutions also don't inspect application traffic (see Figure 1-8). When a connection is established, ZTNA 1.0 trusts that active session implicitly and, therefore, performs no further traffic inspection. If the device is compromised and malware is introduced into the session, there is no means for a ZTNA 1.0 solution to detect any malicious or other compromised traffic and respond accordingly.

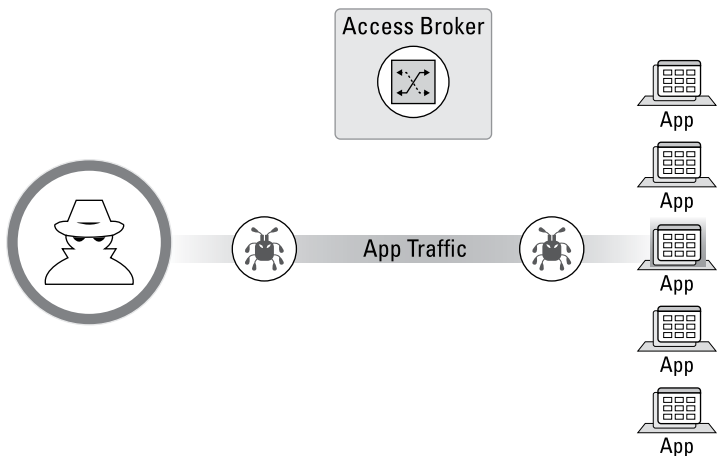


FIGURE 1-8: ZTNA 1.0 provides no security inspection.

Does not protect data

ZTNA 1.0 solutions don't provide data protection — especially the data within private applications (see Figure 1-9). This leaves a good portion of the organization's traffic vulnerable to data exfiltration from malicious insiders or external attackers. Plus, this approach requires additional data loss prevention (DLP) solutions to protect sensitive data in private apps versus SaaS applications. ZTNA 1.0 introduces more complexity and risk because it requires organizations to use multiple-point products to secure data everywhere.

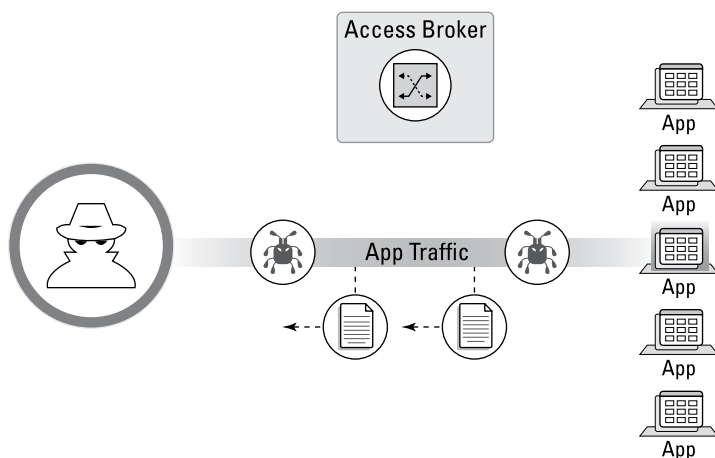


FIGURE 1-9: ZTNA 1.0 provides no data protection.

Does not secure all apps

Finally, ZTNA 1.0 solutions don't provide coverage for all applications (see Figure 1-10). They don't support cloud-based apps or other apps that use dynamic ports or server-initiated applications — like help desk support apps that employ server-initiated connections to remote devices. ZTNA 1.0 solutions don't support SaaS apps either.

Modern cloud-native application stacks are comprised of numerous containers and microservices, often using dynamic IP addresses and port numbers. ZTNA 1.0 access control is completely ineffective in these environments because it requires access to be opened up for broad ranges of IP addresses and ports, effectively defeating the point of Zero Trust.

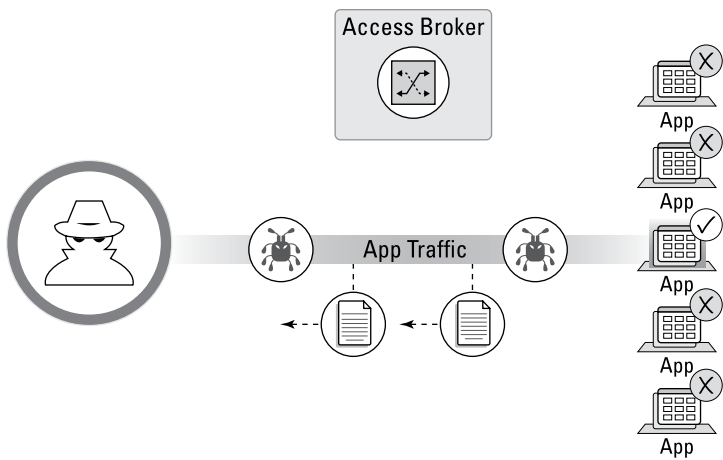


FIGURE 1-10: ZTNA 1.0 can't secure all apps.

As more and more organizations continue on their cloud journey and run their businesses on cloud-native applications, ZTNA 1.0 becomes increasingly obsolete.



TIP

With so many limitations in ZTNA 1.0, you may be wondering, “How did it even make it to the market?” Remember, ZTNA 1.0 was introduced about a decade ago; back then, the world was very different. Prior to ZTNA 1.0, VPN access was really all that was required because all your apps were located in the data center and most users worked in the office. ZTNA 1.0 was introduced to solve some of the problems associated with VPNs as users and apps began to move outside the corporate office and data center. Today, in a world of hybrid network environments and hybrid workforces — in which work is now an activity not a place, and apps and users are everywhere — there is clearly a need for a new approach. Chapter 2 explains how ZTNA 2.0 addresses modern security challenges and goes beyond the limitations of ZTNA 1.0.

IN THIS CHAPTER

- » Implementing least-privilege access
- » Delivering continuous trust verification
- » Enabling continuous security inspection
- » Securing all data
- » Controlling and protecting application access

Chapter 2

Introducing Zero Trust Network Access 2.0

Legacy approaches for secure remote access and out-of-date architectures — like virtual private networks (VPNs) and the initial iteration of Zero Trust Network Access (ZTNA) — are not able to handle the onslaught of new and increasingly sophisticated cyberattacks across our exploding attack surfaces. Clearly, a new approach is needed. This chapter introduces ZTNA 2.0 and explains how it addresses modern security challenges while overcoming the limitations of earlier approaches to enable secure remote access for today's hybrid workforces.

Fully Realizing Least-Privilege Access

ZTNA 2.0 uses stateful application, user, and device identification capabilities to implement least-privilege access (see Figure 2-1).

This means understanding applications from a fundamental perspective at Layer 7 (Application) of the Open Systems Interconnection (OSI) model — beyond low-level network constructs such as Layer 3 (Network [IP address]) and Layer 4 (Transport [port or protocol]) — by continuously gathering information

about Transmission Control Protocol (TCP) session, application handshakes, application behavior, stateful protocols, and more.

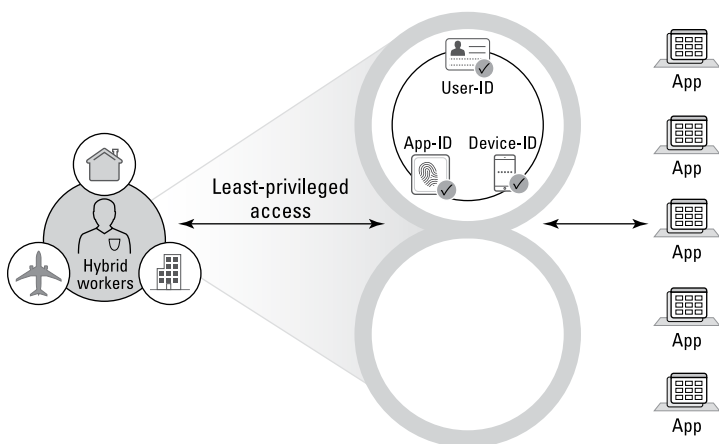


FIGURE 2-1: ZTNA 2.0 uses application, user, and device identification to ensure least-privilege access.

By gaining this level of visibility into applications, especially modern microservices applications, ZTNA 2.0 is able to provide fine-grained controls to prevent exposing sub-app functions or other communication schemas that users do not need access to. At the same time, user and device identification controls continuously gather information about users and their devices. Combining application, user, and device identification moves you beyond simple point-in-time trust assurances (as in ZTNA 1.0) with an environment that provides rich contextual information for making better access control decisions. With ZTNA 2.0, organizations can enable access for any user on any device to the specific application they request and continuously gather additional context to react to changes in real time, dramatically reducing the attack surface while enforcing true least-privilege access.

Enabling Continuous Trust Verification

The core principle of Zero Trust is to remove implicit trust — that is, “Never trust, always verify.” However, without a continuous trust verification capability, the system must assume that

the user, the device, and the app will all behave in a trustworthy manner indefinitely when a connection is established. But a lot can happen to adversely affect trustworthiness after access is granted, including changes in user, device, or application behavior or a security compromise.

Continuous trust verification in ZTNA 2.0 constantly monitors and verifies device posture and any changes to it, along with user and application behaviors, to respond in real time, as needed (see Figure 2-2).

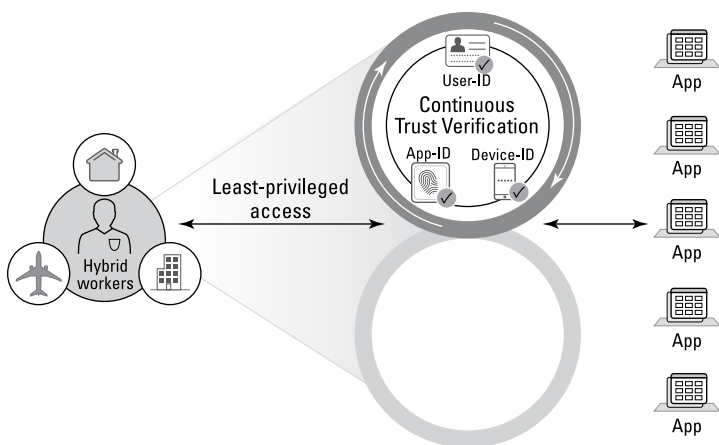


FIGURE 2-2: Continuous trust verification continuously monitors device posture, application behavior, and user behavior even after users gain application access.

Ensuring Continuous Security Inspection

ZTNA 2.0 provides continuous security inspection with threat intelligence, advanced Uniform Resource Locator (URL) filtering, threat prevention, software as a service (SaaS) security, Domain Name System (DNS) security, and more. Deep packet inspection (DPI) and ongoing security inspection capabilities also leverage artificial intelligence (AI)- and machine learning (ML)-powered threat prevention technologies to stop zero-day threats inline (see Figure 2-3).

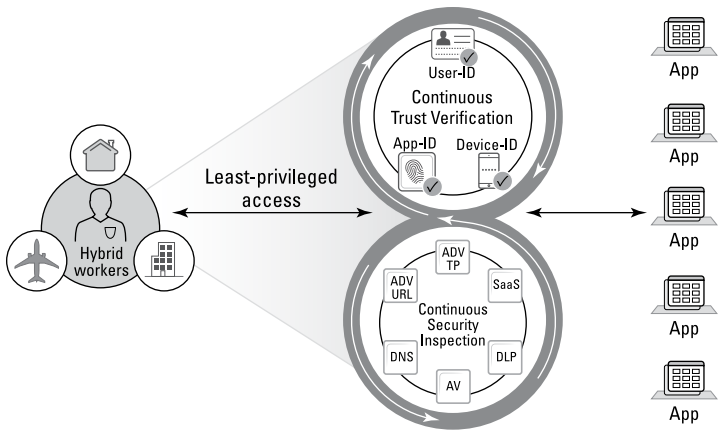


FIGURE 2-3: Continuous security inspection monitors your environment to protect it from threats.

Protecting All Data

ZTNA 2.0 applies advanced data loss prevention (DLP) capabilities consistently to all application data. The same DLP policies are enforced whether the data is in a custom application, a SaaS app, a web app, a public repository, or a database, eliminating the need to guess which apps are protected and what data is secure. Organizations can realize strong data protection and security policies across their apps from a single solution (see Figure 2-4).

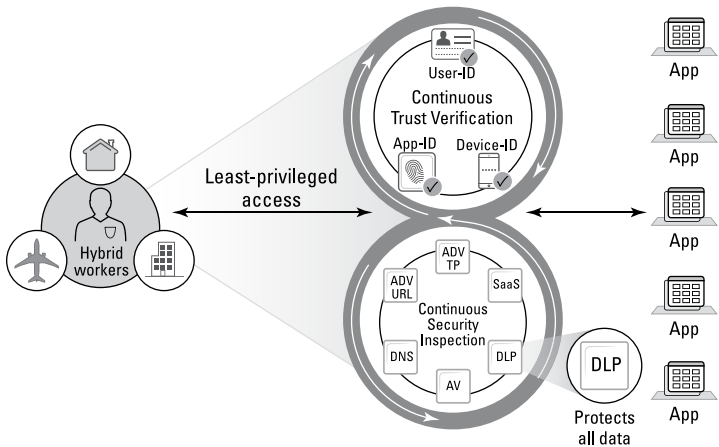


FIGURE 2-4: Consistent data protection applies the same strong data protection and security policies across your environment.

Securing All Apps

ZTNA 2.0 provides consistent security for all applications across your organization. It can be a modern cloud-native microservices-based application that doesn't get restricted by IP addresses and ports, a SaaS app, a custom application, or a legacy application (see Figure 2-5).

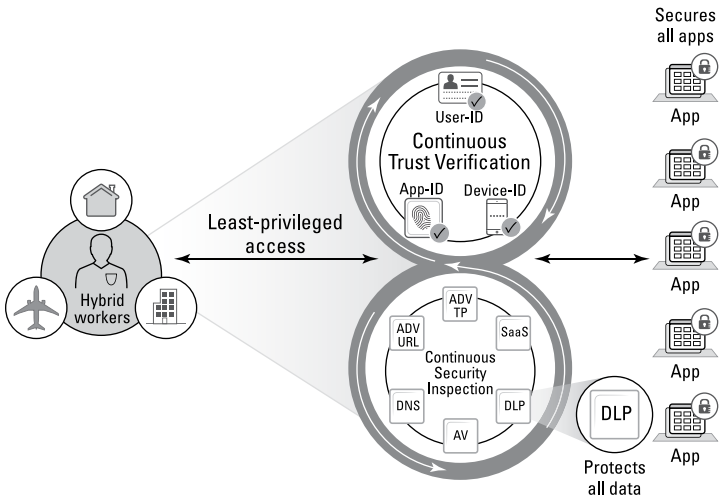


FIGURE 2-5: ZTNA 2.0 provides consistent security for all your applications — whether cloud native, SaaS, custom, or legacy.



REMEMBER

ZTNA 2.0 overcomes the limitations of ZTNA 1.0 solutions and provides better security outcomes to support the digital transformation and hybrid workforce needs facing organizations today. The five key tenets of ZTNA are as follows:

- » **Least privilege:** Uses the most stringent enforcement of the principle of least privilege, providing access control from Layer 3 (Network) through Layer 7 (Application) to dramatically reduce the attack surface.
- » **Continuous trust verification:** When a user's behavior changes, an application's behavior changes, or device posture changes, there has to be a continuous assessment of the trust level granted and the ability to respond appropriately — in real time — to any and all changes.

- » **Continuous security inspection:** All traffic is continuously monitored to protect against all threats — including advanced persistent threats (APTs) as well as zero days — and all threat vectors.
- » **Data protection:** All data is protected with policies applied consistently across all application data, from the data within applications running on legacy mainframes all the way up to the data stored in modern, cloud-native, and collaboration applications.
- » **Consistent security for all apps:** All applications across the organization — including custom applications, cloud-native apps, and SaaS apps — are protected and secured.

- » Recognizing the importance of an exceptional user experience
- » Delivering a simple, unified solution

Chapter **3**

Understanding Critical Capabilities for ZTNA 2.0 Success

This chapter explains why delivering an exceptional user experience and a unified solution is critical to the successful adoption of your ZTNA 2.0 solution.

Delivering an Exceptional User Experience

If you ask your users what they think about your organization's security tools, you probably won't hear "I love the user experience!" Instead, security tools have notoriously been difficult for users to understand and operate, and they typically slow them down. Anti-malware scanning robs your users of valuable memory and makes their computers run slower. Connecting to the virtual private network (VPN) slows down their Internet access and increases latency in their applications. As a result, many users find creative ways to bypass the security controls that are meant to protect them from themselves.

Current ZTNA 1.0 solutions are no different. They rely on physical appliances deployed in colocation facilities that are loosely cobbled together, leveraging the public Internet as their primary backbone. This approach severely limits the reach, scale, and performance of the solution while placing undesired dependency on third-party data centers and suboptimal connections. These solutions also lack true multitenancy to alleviate the challenges of “noisy neighbors” and “fate sharing,” requiring customers to sacrifice security for experience.

To ensure consistent high performance, ZTNA 2.0 solutions should provide a dedicated data plane for each customer, thereby avoiding the “noisy neighbor” problem in ZTNA 1.0 approaches.

ZTNA 2.0 solutions should also be designed with native digital experience monitoring (DEM) capabilities, providing proactive identification of problems, helping resolve issues automatically to reduce the number of trouble tickets IT admins manage, thus providing greater insights and visibility to offer the best experience.



According to Gartner, “Digital experience monitoring (DEM) is a performance analysis discipline that supports the optimization of the operational experience and behavior of a digital agent, human or machine, with the application and service portfolio of enterprises. These users, human or digital, can be a mix of external users outside the firewall and inside it. This discipline also seeks to observe and model the behavior of users as a flow of interactions in the form of a customer journey.”

Providing a Unified Solution

ZTNA 1.0 solutions require you to manage separate policies across different management consoles to completely secure all users and applications. With ZTNA 1.0, it’s impossible to effectively avoid incidents or detect and respond to incidents when management, policies, and data are scattered across your infrastructure.

ZTNA 2.0 solutions provide superior security while delivering uncompromised performance and exceptional user experiences, all in a single unified approach. ZTNA 2.0 provides a truly cloud-native architecture built to secure today’s digital enterprises at cloud scale, providing uncompromised performance backed by

uptime and performance service-level agreements (SLAs) that deliver exceptional user experiences.

Being completely software-based and hardware neutral, ZTNA 2.0 ensures auto-scaling to keep up with changing hybrid workforce and evolving business demands without requiring manual interactions or processes.



REMEMBER

ZTNA 2.0 solutions offer a unified product across all capabilities including ZTNA, SWG, next-generation CASB, FWaaS, DLP, and more.

ZTNA AND SASE

Secure access service edge (SASE) is the convergence of wide-area networking (WAN) and security services in a cloud-delivered services “edge” designed to help organizations modernize their networking and security infrastructures to accommodate the needs of hybrid environments and hybrid workforces.

SASE solutions consolidate multiple point products, including ZTNA, cloud SWG, CASB, FWaaS, and software-defined wide-area networking (SD-WAN), into a single integrated service, reducing network and security complexity while increasing organizational agility.

IN THIS CHAPTER

- » Getting rid of legacy virtual private networks
- » Protecting web applications and Internet traffic
- » Enabling advanced SaaS application protection and data loss prevention

Chapter 4

How to Get Started with ZTNA 2.0

Getting started with Zero Trust Network Access (ZTNA) 2.0 shouldn't be a difficult or overwhelming challenge, and it shouldn't require compromises. It comes down to alignment — mapping your organization's needs to the key concerns or challenges you're facing and solving for those challenges without requiring a massive architectural shift or disruption. This chapter looks at three common use cases that represent some of the biggest challenges organizations are facing today.

VPN Replacement

For years, the standard tool for connecting remote users to a corporate network has been the virtual private network (VPN). VPNs are primarily built to do one thing: Allow remote users to securely access resources inside the corporate network. However, as applications and workloads are increasingly being migrated to the cloud, organizations need more than remote access — they need secure access to cloud applications and the Internet as well.

Legacy VPNs use a hub-and-spoke architecture (see Figure 4-1) to connect remote locations (spokes) to a central office or data center (hub). This location-to-location connectivity 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 4-1: Traditional hub-and-spoke VPN architecture.

The model breaks down when a mixture of cloud and Internet applications are involved. With traditional VPNs, traffic always goes to the corporate VPN concentrator or gateway first, even if the application is hosted in the cloud (see Figure 4-2). As a result, the traffic goes to the VPN gateway at the corporate headquarters or data center and then egresses from the perimeter firewall to the Internet, with the application response going back to headquarters or the data center before it returns to the user. With cloud applications, this traffic essentially follows a “trombone” path, making a lengthy and slow trip to reach an Internet-accessible location. This is sensible from a security perspective, but it doesn’t make sense for network optimization.



TECHNICAL
STUFF

Tromboning refers to the practice of routing network traffic through a control point (such as a firewall). This often involves backhauling traffic, for example, destined to the Internet, across a corporate Multiprotocol Label Switching (MPLS) network and through a central firewall rather than via a more direct route. Tromboning increases network latency and complexity, among other negative impacts.




FIGURE 4-2: Traditional VPNs backhaul traffic to reach the cloud.

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

A key area of focus for many organizations today is replacing outdated VPN technologies that deliver inadequate granularity in controls, poor performance, and an unsatisfactory user experience. VPN replacement initiatives are typically driven by several factors, including the following:

» **Applications moving to a true hybrid model, taking advantage of on-premises, cloud, and multi-cloud environments:**

Legacy VPN technology that trombones or backhauls traffic to an on-premises “concentrator” doesn't scale or deliver the best possible user experience.

» **Changes in enterprise app access requirements:**

Traditionally, employees used managed devices to complete work-related tasks. However, more and more unmanaged devices, like personal phones and tablets, have made their way onto corporate networks and can access corporate applications.

- » Organizations looking for a consistent model of protection and security for all apps universally, not just for web or legacy applications.

VPN technologies were not designed for the rapid scale, high-performance, and consistent delivery of advanced security services required to securely connect a hybrid workforce to the array of applications they require to get their jobs done. Thus, organizations have started replacing outdated VPN deployments with ZTNA solutions.

A number of solutions can address some of these needs, but only ZTNA 2.0 transforms networking and security to support both managed and unmanaged devices while delivering consistent security protection for all applications across the entire organization.

Replacing your VPN with a ZTNA 2.0 solution delivers secure remote access for your branch, home, and mobile workforce to public cloud, private cloud, and data center apps (see Figure 4-3). Key capabilities include continuous trust verification and continuous security inspection that enables the following:

- » Zero Trust model for access to private apps
- » Support for managed or unmanaged client access
- » Consistent protection across the enterprise

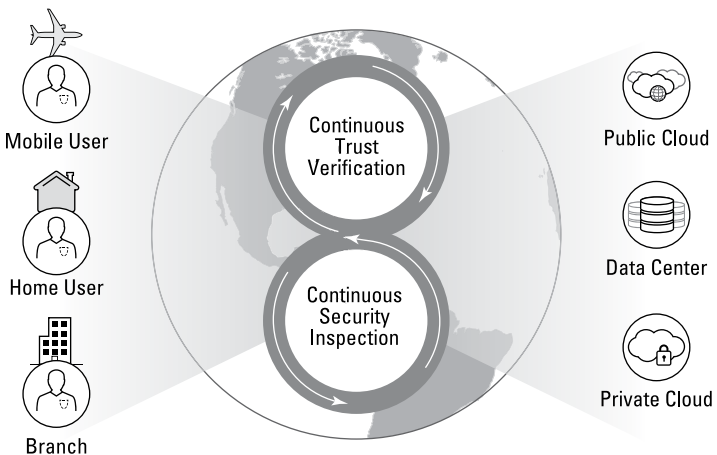


FIGURE 4-3: ZTNA 2.0 for VPN replacement.

Several key benefits of ZTNA 2.0 for VPN replacement projects include the following:

- » Best user experience
- » Unified product
- » Integrated software-defined wide-area network (SD-WAN)



TIP

Replace legacy VPN technologies with a modern ZTNA 2.0 solution that enables secure network access to the remote and hybrid workforce, overcomes performance bottlenecks, and simplifies management.

SECURING PRIVATE ACCESS FOR A FORTUNE 100 CONSULTING FIRM

A Fortune 100 consulting services company was looking for a modern remote access solution that would enable them to retire their aging and unscalable, multi-vendor VPN deployment.

Because of the hodgepodge nature of the VPN solution, they were struggling to achieve consistent visibility and security across the sheer volume of employees and locations around the globe.

Additionally, employee satisfaction with the current mix of solutions was very low. Employees were regularly subjected to slow connects, inconsistent performance, and poor user experiences from site to site and location to location.

Project drivers

- Retire unscalable remote access VPN solution.
- Have consistent visibility and security of employees, regardless of where they are.
- Improve the user experience.

This customer needed a modern replacement for their VPN and chose the Palo Alto Networks ZTNA 2.0 solution. With ZTNA 2.0, they're now able to connect all 350,000 users across 158 countries consistently, while also providing secure direct-to-Internet connectivity for the

(continued)

(continued)

hundreds of branch offices across the globe. What's more, ZTNA 2.0 ensures consistent and secure access to all apps, including legacy apps, across 30-plus data centers and cloud locations.

Impact

- 350,000 users secured across 158 countries
- Local Internet with security from the cloud protecting hundreds of offices globally
- ZTNA for thousands of apps across 30-plus data centers and cloud locations

Securing Internet Access

Organizations use many applications, some located on-premises and others in the cloud. As enterprises and their mobile and hybrid workforces grow, it becomes increasingly difficult to protect remote users from threats as they access these various applications.

On-premises applications are typically accessed through a remote access VPN. However, when users access Internet-based applications and services, they're disconnected from the VPN and exposed to risk. Organizations use secure web gateways (SWG) to provide secure Internet access when remote users are disconnected from the VPN.

An SWG typically acts as a *proxy* (intermediary) between users and Internet resources to protect users from web-based threats in addition to applying and enforcing corporate acceptable use policies. Instead of connecting directly to a website, a user is directed to the SWG, which is then responsible for connecting the user to the desired website and performing functions such as Uniform Resource Locator (URL) filtering, web visibility, malicious content inspection, web access controls, and other security measures.



REMEMBER

SWGs enable companies to:

- » Block access to inappropriate websites or content based on acceptable use policies

- » Enforce security policies to make Internet access safer
- » Protect data against unauthorized transfer

However, legacy SWGs are typically deployed as appliances on corporate networks, requiring user traffic to be backhauled to the SWG — which is often located in a corporate data center. This inefficient routing of traffic increases latency and negatively impacts the user experience.

Another challenge of legacy SWGs is that they're typically stand-alone solutions that lack the ability to coordinate workflows, reporting, or logging with other security infrastructure in the organization. This can lead to increased complexity over time because organizations often have multiple security point products that make their security operations less efficient and less effective.

As organizations look for ways to improve their employee experience when accessing the Internet and web applications, the cloud SWG capabilities in ZTNA 2.0 offer an effective solution that removes latency and improves overall security posture (see Figure 4-4).

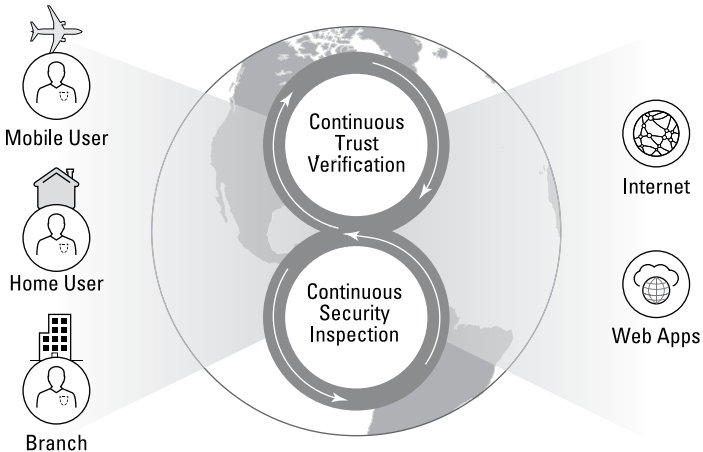


FIGURE 4-4: ZTNA 2.0 for SWG replacement/cloud SWG.



TIP

Some key requirements to look for when evaluating ZTNA 2.0 as a replacement for legacy SWG products include the following:

- » **Doesn't require significant network changes:** Organizations want to be able to simply maintain existing proxy-based approaches to minimize disruptions and major network overhauls.
- » **Offers an optional agent-based approach:** Having the option for an agent to be installed on user endpoints is desirable, but it shouldn't be the only deployment model available.
- » **Delivers consistent policy enforcement:** The solution must deliver consistent enforcement across a hybrid workforce including mobile, home, and branch users.

SECURING INTERNET ACCESS FOR A FORTUNE 100 PHARMACEUTICAL COMPANY

A Fortune 100 pharmaceutical company wanted to reduce its multivendor, on-premises hardware deployments and modernize its infrastructure with cloud-delivered security.

Because more of the tools and apps employees needed to do their work were migrating to the cloud, existing solutions couldn't offer a seamless experience and match expectations, which caused poor overall user satisfaction with the current solutions.

This customer needed a modern, cloud-delivered security approach and chose the Palo Alto Networks ZTNA 2.0 solution. They were able to easily migrate all 100,000 users within three months without re-architecting their network, leveraging the explicit proxy capabilities of ZTNA 2.0.

Their new cloud-native solution consolidated and eliminated their on-premises proxy hardware and allowed them to realize an improved security posture across all users and locations. They also deployed the native autonomous digital experience management (ADEM) capabilities of Prisma Access to ensure exceptional user experiences for all hybrid workers.

Project drivers

- Migrating to cloud
- Reducing on-premises hardware
- Improving user experience

Impact

- 100,000 users migrated in less than three months
- On-premises SWG hardware eliminated by cloud-native solution
- Dramatically improved security posture
- Exceptional user experience with ADEM

Advanced SaaS Security

Years ago, companies typically kept all their applications and data in an on-premises data center. In this environment, companies had complete visibility into and granular control over who was accessing their applications and data and when, as well as which devices (typically desktop or laptop computers) were being used to access them.

Over time, as companies moved their data to the cloud and began using cloud services such as SaaS applications, they discovered that they no longer had insight into who was accessing and using their cloud applications and data or — thanks to the advent of mobile technologies such as laptops and smartphones — the devices being used to access these cloud services. In addition, the ubiquity and ease of adoption of SaaS applications often leads to “shadow IT,” where users leverage unsanctioned or unapproved applications for business purposes that inadvertently exposes sensitive data to increased risk.

This lack of visibility makes it difficult for companies to protect their data and opens them up to a host of enterprise security risks, such as data breaches, regulatory noncompliance, malware, ransomware, and more.

To address these challenges, security vendors developed cloud access security broker (CASB) solutions. CASBs help organizations

discover where their data is located across SaaS applications, and when it's in motion across cloud services environments, on-premises data centers, and mobile workers. A CASB also enforces an organization's security, governance, and compliance policies, allowing authorized users to access and consume cloud applications while enabling organizations to protect their sensitive data effectively and consistently across multiple locations.

However, conventional CASB solutions are unable to quickly onboard new cloud applications because they rely on static application libraries that are manually populated. Modern collaboration apps like Slack, Zoom, Confluence, Jira, and others, where users spend most of their time today, are typically not covered by the application programming interface (API) protections offered in these CASB solutions.

A conventional CASB solution offers basic cloud security capabilities that are limited in breadth and depth, providing only piecemeal security. For example, its data loss prevention (DLP) capabilities are quite basic and inaccurate, covering only data in certain SaaS apps while completely detached from any DLP protecting the rest of the enterprise. They also lack the essential threat protection mechanisms that detect endless threat variations that cybercriminals constantly utilize to target SaaS applications.



TECHNICAL
STUFF

When CASB solutions were first developed, they were designed as a stand-alone proxy-based point solution. The issue with proxy-based CASBs is that they require complex traffic redirection from the network firewall with proxy auto-configuration (PAC) agents and log collectors, causing significant architectural and operational complexity together with high cost of ownership.

Enterprises using legacy CASB solutions today can't keep up with the rapid growth of SaaS applications and shadow IT, the ubiquitous growth of data, or the increasing numbers of hybrid and remote workers. Replacing legacy CASB with next-generation CASB capabilities delivered in a secure access service edge (SASE) architecture that encompasses ZTNA 2.0, enables enterprises to safely embrace cloud services with continuous trust verification and security inspection, and capabilities that include the following (see Figure 4-5):

- » SaaS app visibility and control
- » Protection of sanctioned SaaS apps
- » Advanced DLP

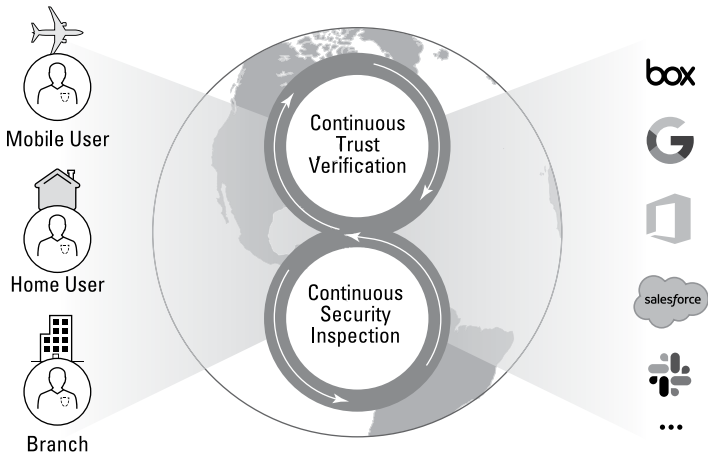


FIGURE 4-5: ZTNA 2.0 for advanced SaaS app security/next-generation CASB.

ADVANCED SaaS APP SECURITY FOR A LARGE AUTOMOTIVE SUPPLIER

A global automotive technology leader, with more than 190,000 people across hundreds of manufacturing facilities, dozens of major technical centers worldwide, and a presence in more than 40 countries, was relying on more apps in the cloud. They needed better visibility and granular control of known and unknown SaaS applications, consolidated management of multivendor products, and inline threat inspection.

The company was also looking for simple policy creation and deployment without leveraging a proxy or agents and wanted to eliminate the need to synchronize risks, policies, and goals across a separate layer of the stack.

(continued)

(continued)

The Palo Alto Networks ZTNA 2.0 solution with next-generation CASB capabilities enabled the company to eliminate the requirement to update and configure agents for inline inspection and protect unmanaged endpoints.

Project drivers

- Massive cloud/SaaS adoption
- Visibility and control over known and unknown apps
- Reduce complexity, consolidate security

Impact

- Cloud-delivered security simplified deployment and policy creation
- Dramatically increased visibility and control of all apps
- Consistent protections for 190,000 users globally

IN THIS CHAPTER

- » Ensuring complete visibility and control
- » Delivering continuous trust verification
- » Securing all apps in a unified solution
- » Applying complete security inspection
- » Preventing data loss across all environments
- » Guaranteeing application uptime and performance
- » Reducing complexity and cost in a single solution

Chapter **5**

Ten (or So) Questions to Ask Your ZTNA 2.0 Vendor

Here are some important questions to help you evaluate potential Zero Trust Network Access (ZTNA) 2.0 vendors and their solutions.

Do You Provide Full Layer 7 Application Visibility?

Users increasingly leverage a variety of applications — including a host of software as a service (SaaS) applications from a multitude of devices and locations — for work-related as well as personal purposes. Many applications, such as instant messaging (IM), peer-to-peer (P2P) file sharing, and Voice over Internet Protocol

(VoIP), are capable of operating on nonstandard or dynamic ports and IP addresses.

Furthermore, 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 apps (sanctioned, tolerated, unsanctioned). Thus, a ZTNA solution that identifies applications based on arbitrary Layer 3 port assignments and is limited to Layer 3 or 4 access control is no longer sufficient to protect your enterprise.



TIP

Look for a ZTNA 2.0 solution that can classify traffic by application on all ports, all the time, by default — and doesn't create an administrative burden by requiring you to research which applications use which ports in order to configure appropriate policies and rules. A complete ZTNA 2.0 solution provides complete Layer 7 (Application) visibility into application usage along with capabilities to understand and control their use.

Do You Provide Continuous Trust Verification?

The basic principle of Zero Trust is “Never trust, always verify” — not “Never trust, verify once” or “Never trust, verify occasionally.” Trusting an entity based on static account credentials vetted once on a device that appears legitimate at one point in time is a recipe for disaster. Cybercriminals leverage this flawed inherent trust model to move freely within a network environment after they've breached the perimeter defenses.



TIP

A robust ZTNA 2.0 solution provides continuous trust verification based on individual user behavior leveraging machine learning (ML) to determine risk and identify potential threats. Access to the network or an application should be allowed only after carefully vetting of users and devices that includes multifactor authentication (MFA). But it doesn't stop there. Ongoing trust verification should occur continuously and seamlessly throughout the session to ensure that the security posture of the user or device hasn't changed or been compromised.

Do You Consistently Secure All Apps in a Single Product?

As discussed in Chapter 1, point security solutions that only protect specific applications or support limited use-case scenarios lead to complexity, inefficiency, and, ultimately, a weaker security posture. Users will find creative new ways to bypass security controls that are confusing and inconvenient to use. Security teams are more prone to mistakes in configuring and operating tools with different operating systems, interfaces, and syntaxes, and they'll be overwhelmed with alerts that can't be easily correlated with specific threats in an integrated solution.



TIP

Your ZTNA 2.0 solution should consistently secure all your applications, including legacy custom apps, modern microservices-based cloud-native apps, SaaS apps, and more in a single unified product.

Do You Apply Complete Security Inspection?

A ZTNA 2.0 solution must do more than simply allow or block traffic based on limited inspection of packet headers and enforcement of static firewall rules. It must also prevent advanced malware, including ransomware, and inspect for known and unknown threats in both encrypted and unencrypted application and data traffic — again, across all application traffic, not just private app traffic.



TIP

A complete ZTNA 2.0 solution must provide complete security inspection and control including malware and threat prevention.

Do You Consistently Protect All Enterprise Data?

Consistent data protection requires consolidating data protection policies across all environments and every data communication vector. Disjointed data protection policies and configurations for

different SaaS applications, on-premises repositories, email communications, local storage, and so on, cause security blind spots, management complexity, inconsistent controls, and shadow IT.



TIP

Choose a ZTNA 2.0 solution that enables a consistent data loss prevention (DLP) policy across every environment where data lives and flows, regardless of its location.

Do You Provide Uptime and Performance SLAs for All Apps?

Security tools that negatively impact application uptime and performance will contribute to a poor user experience that will ultimately lead to more shadow IT as users look for new ways to circumvent the tools that are meant to protect them.



TIP

Modern ZTNA 2.0 solutions are cloud-delivered and must, therefore, provide reliability and performance guarantees that ensure your organization's users can use the applications they require — including internal and SaaS-based — securely and efficiently, when they need them. Ensure your ZTNA 2.0 vendor provides uptime and performance service-level agreements (SLAs) that meet your organization's needs.

Do You Have a Single Unified Product to Secure the Enterprise?

Siloed security tools that can't be easily integrated with other solutions add cost and complexity to your environment and can delay critical threat detection, correlation, identification, and response. This added complexity ultimately leads to increased management overhead while significantly increasing risks and exposure.



TIP

A ZTNA 2.0 vendor should offer a single unified solution, like a secure access service edge (SASE), that protects your entire enterprise — including users, applications, devices, and data — regardless of their location to reduce risk and deliver better security outcomes.

Glossary

ADEM: *See* Autonomous Digital Experience Management (ADEM).

AI: *See* artificial intelligence (AI).

antivirus (AV): Software that is designed to detect and prevent computer viruses and other malware from infecting a system.
See also malware.

API: *See* application programming interface (API).

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

artificial intelligence (AI): The ability of a computer to interact with and learn from its environment and to automatically perform actions without being explicitly programmed.

Autonomous Digital Experience Management (ADEM): A feature in Palo Alto Networks Prisma Access that provides SASE-native digital experience monitoring and complete visibility to autonomously remediate user connectivity issues before or when they arise.
See also secure access service edge (SASE).

AV: *See* antivirus (AV).

C2: *See* command-and-control (C2).

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

data loss prevention (DLP): An application or device used to detect the unauthorized storage or transmission of sensitive data.

deep packet inspection (DPI): An advanced method of examining and managing network traffic that extends beyond the initial packet headers.

DLP: See data loss prevention (DLP).

DNS: See Domain Name System (DNS).

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. See *also* Internet Protocol (IP).

DPI: See deep packet inspection (DPI).

EDR: See endpoint detection and response (EDR).

endpoint detection and response (EDR): A category of tools used to detect and investigate threats on endpoints. EDR tools typically provide detection, investigation, threat hunting, and response capabilities.

endpoint protection platform (EPP): An integrated suite of endpoint security technologies such as antivirus, data encryption, data loss prevention, personal firewall, and port and device control.

EPP: See endpoint protection platform (EPP).

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 attack.

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

FWaaS: See firewall as a service (FWaaS).

hybrid cloud: An environment consisting of resources from multiple public and/or private clouds that provide application and data portability across clouds. See *also* private cloud *and* public cloud.

IM: See instant messaging (IM).

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

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

intrusion prevention system (IPS): A hardware or software application that both detects and blocks suspected network or host intrusions.

IP: *See* Internet Protocol (IP).

IPS: *See* intrusion prevention system (IPS).

LAN: *See* local-area network (LAN).

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

machine learning (ML): A method of data analysis that enables computers to analyze a data set and automatically perform actions based on the results without being explicitly programmed.

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

MFA: *See* multifactor authentication (MFA).

ML: *See* machine learning (ML).

MPLS: *See* Multiprotocol Label Switching (MPLS).

multi-cloud: An environment consisting of resources from multiple public and/or private clouds, but that does not necessarily provide application and data portability across clouds (that is, the different cloud environments may operate as siloed clouds). It's important to note that although all hybrid cloud environments are also multi-cloud environments, not all multi-cloud environments are hybrid cloud environments. *See also* hybrid cloud, private cloud, *and* public cloud.

multifactor authentication (MFA): An authentication mechanism that requires two or more of the following factors: something you know, something you have, or something you are. For example, a user may authenticate with their username and password (something they know) and a one-time passcode sent to a mobile phone that has previously been registered with the organization (something they have).

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.

network traffic analysis (NTA): A category of tools used to intercept, record, and analyze network traffic patterns to detect and respond to anomalies and suspicious activities using a combination of machine learning, behavioral modeling, and rule-based detection. *See also* machine learning.

NTA: See network traffic analysis (NTA).

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

OSI model: See Open Systems Interconnection (OSI) model.

P2P: See peer-to-peer (P2P).

PAC: See proxy auto-configuration (PAC) file.

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

proxy auto-configuration (PAC) file: A web-based rule set written in JavaScript that advises your endpoint on how to direct its traffic for a given URL: either via a web proxy or directly to the Internet. It can contain information including the IP address of the website, the IP address of the user, and the host that requested the website. See also Uniform Resource Locator (URL).

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

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

RDP: See Remote Desktop Protocol (RDP).

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

SaaS: See software as a service (SaaS).

SDP: See software-defined perimeter (SDP).

SD-WAN: See software-defined wide-area network (SD-WAN).

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

secure web gateway (SWG): A security platform or service that is designed to maintain visibility into all types of traffic, while stopping evasions that can mask threats. Additional functionality may include web content filtering and credential theft prevention.

security information and event management (SIEM): A system that provides real-time collection, analysis, correlation, and presentation of security logs and alerts. Security operations center (SOC) analysts use SIEM tools to manage security incidents, and detect and respond to potential threats quickly. *See also* security operations center (SOC).

security operations center (SOC): A facility that provides cybersecurity monitoring, assessment, defense, and remediation for enterprise compute and network resources, including on-premises and cloud environments.

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

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

SIEM: *See* security information and event management (SIEM).

SLA: *See* service-level agreement (SLA).

SOC: *See* security operations center (SOC).

software as a service (SaaS): A cloud-based software distribution model in which a third-party provider hosts applications that it makes available to customers over the Internet. The software vendor hosts and maintains the servers, databases, and code that constitute an application.

software-defined perimeter (SDP): A software-defined perimeter secures all connections to services running on a network infrastructure at all layers, based on the level of security you define and establish.

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. *See also* wide-area network (WAN).

SSH: *See* Secure Shell (SSH).

SWG: *See* secure web gateway (SWG).

tactics, techniques, and procedures (TTPs): The behaviors, methods, strategies, and tools used by threat actors to attack a target.

TCP: *See* Transmission Control Protocol (TCP).

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.

tromboning: The practice of routing network traffic through a control point (such as a firewall).

TTPs: *See* tactics, techniques, and procedures (TTPs).

UDP: *See* User Datagram Protocol (UDP).

UEBA: *See* user and entity behavior analytics (UEBA).

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

URL: *See* Uniform Resource Locator (URL).

user and entity behavior analytics (UEBA): A type of cybersecurity solution or feature that discovers threats by identifying activity that deviates from a normal baseline. Although UEBA can be used for a variety of reasons, it’s most commonly used to monitor and detect unusual traffic patterns, unauthorized data access and movement, or suspicious or malicious activity on a computer network or endpoints.

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

virtual private network (VPN): A VPN creates a private connection, known as a *tunnel*, to the Internet. All information traveling from a device connected to a VPN will get encrypted and go through this tunnel. When connected to a VPN, a device will behave as if it’s on the same local network as the VPN. The VPN will forward device traffic to and from the intended website or network through its secure connection.

Voice over Internet Protocol (VoIP): Telephony protocols that are designed to transport voice communications over TCP/IP networks. *See also* Transmission Control Protocol (TCP) *and* Internet Protocol (IP).

VoIP: *See* Voice over Internet Protocol (VoIP).

VPN: *See* virtual private network (VPN).

WAN: *See* wide-area network (WAN).

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

Zero Trust: Zero Trust is a strategic initiative that helps prevent successful data breaches by eliminating trust from your organization. Rooted in the principle of “never trust, always verify,” Zero Trust is designed to prevent lateral movement.

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.

ZTNA: See Zero Trust Network Access (ZTNA).

Notes

Notes

Notes

Notes

Notes

Notes

Zero Trust with Zero Exceptions

**ZTNA 1.0 is over. Secure the future of hybrid work with ZTNA 2.0.
Only available with Prisma® Access.**

Palo Alto Networks Prisma® Access protects the hybrid workforce with the superior security of ZTNA 2.0 while providing exceptional user experiences from a simple, unified security product. Purpose-built in the cloud to secure at cloud scale, only Prisma Access protects all application traffic with best-in-class capabilities while securing both access and data to dramatically reduce the risk of a data breach.

Learn how Prisma Access secures today's hybrid workforce without compromising performance, backed by industry-leading SLAs to ensure exceptional user experiences.

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



Get Started on ZTNA Today

The hybrid workforce and direct-to-app architectures have rendered traditional security solutions obsolete while exponentially increasing the attack surface. At the same time, threats are increasing in frequency and sophistication while the proliferation of disparate security tools creates operational complexity. Existing cloud-based security solutions provide too much access with too little protection, deliver inconsistent and incomplete security across applications, and offer poor performance and user experiences. Zero Trust Network Access 2.0 offers a better path forward.

Inside...

- Discover use cases to get started on your ZTNA 2.0 journey
- Understand the differences between ZTNA 2.0 and legacy ZTNA solutions
- Learn the five tenets of ZTNA 2.0
- Identify questions to ask your ZTNA vendor
- See how a unified solution offers exceptional user experience



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 200 *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: 978-1-119-98379-8

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.