

Brought to you by:



Software Firewalls

for
dummies[®]
A Wiley Brand



Extend Zero Trust
to cloud applications

Discover software
firewall types

Explore software
firewall use cases

Palo Alto Networks
Special Edition

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 and 2022), Comparably Best Companies for Diversity (2021), and HRC Best Places for LGBTQ Equality (2022). For more information, visit www.paloaltonetworks.com.



Software Firewalls

Palo Alto Networks Special Edition

by Lawrence Miller

for
dummies[®]
A Wiley Brand

Software Firewalls For Dummies® , Palo Alto Networks Special Edition

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

Copyright © 2023 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 978-1-394-16555-1 (pbk); ISBN 978-1-394-16556-8 (ebk)

Publisher's Acknowledgments

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

Editor: Elizabeth Kuball

Acquisitions Editor: Traci Martin

Editorial Manager: Rev Mengle

Client Account Manager:

Cynthia Tweed

Production Editor:

Tamilmani Varadharaj

Table of Contents

INTRODUCTION	1
About This Book	1
Foolish Assumptions	2
Icons Used in This Book	3
Where to Go from Here	3
CHAPTER 1: Recognizing Current Trends and Challenges	5
Growth of Hybrid Clouds and Application Portability	5
Security Challenges in Cloud and Virtualized Environments	8
The evolving and modern threat landscape	9
A tale of two teams	10
Security of the cloud versus security in the cloud	12
Compliance requirements	15
CHAPTER 2: Understanding Software Firewalls and Zero Trust	17
What Is Zero Trust?	17
How Do Software Firewalls Help You Achieve Zero Trust?	20
Extending Zero Trust Practice to Cloud Applications	21
CHAPTER 3: Discovering Software Firewall Types	23
Virtual Firewalls	23
Cloud Firewalls	26
Container Firewalls	28
Cloud-Delivered Security Services	31
CHAPTER 4: Exploring Software Firewall Use Cases	33
Public Cloud	33
Application security detects hard-to-find threats	34
Outbound traffic protection stops exfiltration	34
Filtering and inspection boosts developer security	35
Private Cloud	36
Segmentation and micro-segmentation protect against lateral movement	37
Augment software-defined networking with threat prevention	38
VDI security addresses threats to remote and distributed workforces	38

Hybrid Cloud and Multi-Cloud	39
Virtualized Branch	40
Achieving compliance with local branch segmentation	40
Software-based perimeter security simplifies deployment.....	41
Secure SD-WAN increases performance and network return on investment.....	42
5G	43
CHAPTER 5: Ten Questions to Ask Your Software Firewall Vendor	45
Do They Stop Zero-Day Threats?	45
Do They Provide Least-Privilege Access Control?	46
Can Security Be Consolidated into a Security Platform?	46
Can They Provide Consistent and Future-Proof Protection?	47
Can You Dial Up/Down Security with Flexible Consumption?	47
Do They Provide Single-Pane-of-Glass Management?	48
Can They Secure Any Cloud and Application Architecture Model?	48
Do They Work with Your Automation/Orchestration Tools?	50
Are They Proven to Accelerate Security Posture?	50
Do They Have a Track Record of Delivering High ROI?	51
GLOSSARY	53

Introduction

As more enterprises adopt or expand their cloud strategies, they must address a fundamental security dilemma: Physical next-generation firewalls are not cloud-friendly, and basic cloud service provider (CSP) firewalls are not effective against modern cybersecurity attacks. Network firewalls have traditionally been configured on hardware that is specifically designed to handle the high throughput and intense processing requirements of enterprise firewalls — particularly next-generation firewalls.

Fortunately, technology has evolved, and the key to innovation today lies in software. Practically everything today is “software-defined,” and, in the words of Microsoft CEO Satya Nadella (among others), “every company is now a software company.” Software firewalls have emerged as an optimal solution for a variety of cloud and virtualized use cases including private clouds, software-defined networks (SDNs), multi-cloud, hybrid, edge, containers, and 5G.

About This Book

Software Firewalls For Dummies, Palo Alto Networks Special Edition, consists of five chapters that explore the following:

- » The modern threat landscape, security challenges, and limitations of existing technologies and approaches in hybrid and multi-cloud environments (Chapter 1)
- » Zero Trust and how to extend Zero Trust to the cloud with software firewalls (Chapter 2)
- » Different types of software firewalls (Chapter 3)
- » Software firewall use cases (Chapter 4)
- » Key questions to ask your software firewall 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 cloud operations professional, security professional, or decision-maker looking for a better way to secure your organization's application workloads in one or many flavors of cloud and virtualized environments. Perhaps you're a chief information security officer (CISO), focused on trying to stay ahead of the curve and making sure that everything fits together, solutions are implementable, and the business keeps moving. Or you may be a cloud infrastructure architect, engineer, or application developer focused on ensuring your organization's complex multi-cloud platform — especially its architecture, infrastructure, and applications — meets the needs of your dynamic business. Or perhaps you're a network security architect or engineer looking for the best security solutions to consistently protect critical business applications and data, while gaining full end-to-end visibility across the infrastructure.

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, after reading it, you'll be well aware of the advantages of software firewalls in modern cloud and virtualized environments.

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.



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.

Where to Go from Here

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 <https://paloaltonetworks.com>.

- » Adopting a hybrid cloud strategy
- » Looking at security challenges in cloud and virtualized environments and the evolving threat landscape

Chapter 1

Recognizing Current Trends and Challenges

In this chapter, you explore the growth of hybrid clouds and application portability and how cloud-first/cloud-native strategies have become key business enablers. You also learn about unique security challenges in cloud and virtualized environments, as well as the evolving and modern threat landscape including ransomware as a service (RaaS) and morphing threats that are increasingly difficult to detect.

Growth of Hybrid Clouds and Application Portability

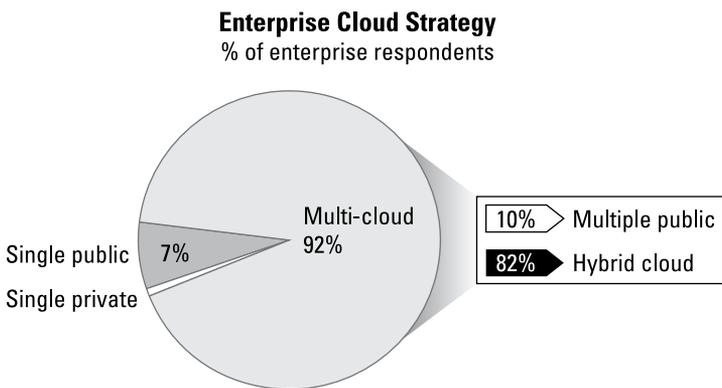
Modern enterprises are revolutionizing and reinventing the way they do business to stay market competitive with a focus on the following:

- » **Focusing on innovation** that grows the business by opening up and freeing up resources

- » **Delivering superior customer and employee experiences** that please and make them more productive
- » **Enabling agility and scalability** of products, processes, and the overall business to stay ahead of the competition

Modern enterprises have extended their private cloud and virtualized environments to support hybrid cloud and multi-cloud strategies, all of which are critical to these business initiatives.

According to the *Flexera 2021 State of the Cloud Report*, 92 percent of enterprises today have a multi-cloud strategy and 82 percent have a hybrid cloud strategy (see Figure 1-1).



Source: Flexera 2021 State of the Cloud Report

FIGURE 1-1: Enterprises have embraced multi-cloud and hybrid cloud strategies.



**TECHNICAL
STUFF**

All hybrid clouds are multi-clouds, but not all multi-clouds are hybrid clouds. A hybrid cloud is comprised of virtualized (for example, VMware, Nutanix, Linux Kernel-based Virtual Machine [KVM], and so on) compute, storage, and networking that combines on-premises infrastructure — or a private cloud — with a public cloud (such as Amazon Web Services [AWS], Google Cloud Platform [GCP], Microsoft Azure, Alibaba, and so on). Multi-cloud consists of two or more public clouds. Both hybrid clouds and multi-clouds allow data and applications to move between the cloud environments.

In the not-too-distant past — before cloud computing was everywhere — business applications were hosted in centralized enterprise data centers. As software-as-a-service (SaaS) applications became increasingly popular, network traffic for remote

workers and branch users had to be backhauled (also referred to as “hairpinning”) through the data center, which had network firewalls deployed at the edge of its perimeter for centralized protection.

This centralized model no longer aligns with modern business needs for innovation, agility, and great customer and employee experiences. As businesses increasingly move to hybrid cloud and multi-cloud environments to achieve greater flexibility, this problem has only been exacerbated. As a result, the enterprise attack surface has expanded by becoming a complex, hyperconnected environment composed of multiple clouds, which are increasingly difficult to protect (see Figure 1-2).

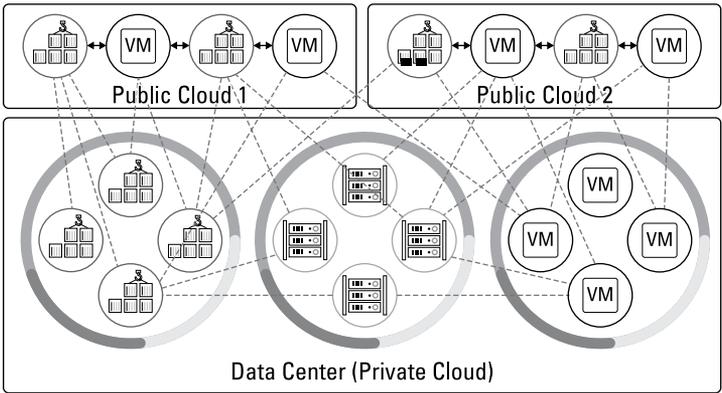


FIGURE 1-2: Today’s hyperconnected multi-cloud, hybrid environments.

For some organizations, the cloud migration journey often begins with adopting external SaaS applications, either as a business strategy or as a result of shadow IT — or both. They then take the next logical step toward virtualization of data center infrastructure, leveraging hypervisor and software-defined networking technologies, and then lifting-and-shifting traditional client/server applications onto virtual machines (VMs) and virtual desktop infrastructure (VDI).

Many organizations are rapidly adopting public infrastructure-as-a-service (IaaS) solutions in developing their modern applications as part of a deliberate application strategy, then lifting-and-shifting, building cloud-native, or both, in migrating their on-premises applications and workloads to the public cloud.

Finally, many organizations take advantage of platform-as-a-service (PaaS) solutions and other cloud-native innovations with serverless technology and containerization as some of the dominant modern application architectures. As organizations undertake this journey, they quickly find that on-premises hardware firewalls can't be deployed in their virtualized data centers, private, and public cloud environments. This is where software firewalls come into the picture.



REMEMBER

Shadow IT refers to the use of technology solutions without the knowledge or approval of the IT department. Shadow IT increases enterprise security and compliance risk, among many other management challenges.



TIP

According to the *ESG Research Report: Cloud-native Applications*, 88 percent of organizations currently deploy production applications and workloads on public cloud infrastructure services.

Security Challenges in Cloud and Virtualized Environments

Modern enterprises adopt cloud-first and cloud-native strategies to take advantage of agility, flexibility, innovation, and cost savings, among other benefits. However, the cloud and virtualized environments also create new challenges and risks that businesses must address, such as:

- » **Incomplete visibility:** Ensuring end-to-end visibility is challenging in any environment, but even more so for virtualized and cloud environments. Traffic between VMs on the same physical host typically only traverses the hypervisor (software creating and running VMs). East-west traffic between VMs in a virtualized data center or virtual private clouds (VPCs) in cloud environments typically never cross a firewall (deployed to inspect and control network traffic).
- » **Limited segmentation:** Logical segmentation in physical networks is typically achieved with virtual local area networks (VLANs) configured on network switches. VLANs are less often configured in virtualized and cloud environments. Instead, segmentation is often achieved by configuring separate virtual (or virtual private) networks, but backplane

connections and hypervisor traffic often remain open and flat across the virtual layer.

- » **Inadequate protection:** Basic firewalls and network security group-like tools only provide basic port/source/protocol filtering capabilities. They can't inspect the traffic itself, allowing for many loopholes for malicious traffic to evade these basic capabilities.
- » **Speed of DevOps:** Application development teams are under constant pressure to deliver new software as quickly as possible. Moving applications into production has been extensively automated within continuous integration/continuous deployment (CI/CD) pipelines and infrastructure as code (IaC) to rapidly build up and tear down application infrastructure at scale. As a result, security often struggles to keep up with the rapid pace of business innovation.

Let's take a closer look at some of the security challenges in virtualized and cloud environments that create attack opportunities for threat actors.

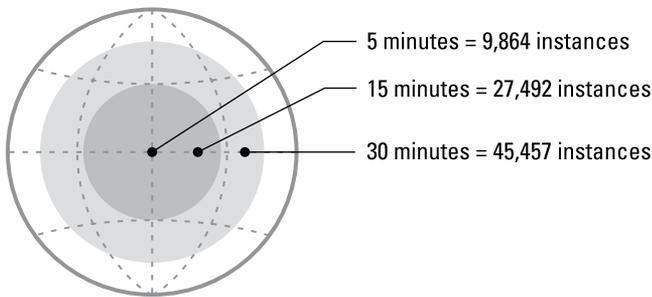
The evolving and modern threat landscape

Many organizations find cloud security more difficult than traditional on-premises security due to a number of factors, including the following:

- » **Expanding attack surface:** The trend toward applications being hosted across a wide range of environments — in addition to remote work and mobile user access — creates a significantly larger attack surface. Additionally, the growth of the integrated supply chain applications increases risk when vendors and partners fail to maintain an adequate security posture within their own environments.
- » **More sophisticated threats:** Today's modern threats use advanced techniques to exploit vulnerabilities and evade detection. These techniques include, for example, encrypting malware, establishing command-and-control (C2) communications, exploiting zero-day vulnerabilities, and leveraging constantly evolving — including artificial intelligence (AI)-enabled — malware variants to avoid detection. Additionally, the rise of RaaS has made it easier for

practically any threat actor to carry out sophisticated attacks, regardless of their skill level. As a result, ransomware (and data exfiltration) attacks are now among the top cloud security breaches.

» **Shorter action windows:** Mean-time-to-detect (MTTD) and mean-time-to-respond (MTTR) are typically measured in days — the average time to identify and contain a data breach is 277 days according to the IBM Security *Cost of a Data Breach Report 2022* — but ransomware can begin to encrypt your data in just minutes after penetrating the network and gaining access to data sources. In a recent Palo Alto Networks Unit 42 internal study, an advanced threat replicated itself into more than 45,000 instances in just 30 minutes (see Figure 1-3).



Source: Palo Alto Network Unit 42 internal study

FIGURE 1-3: Sophisticated threats proliferate rapidly across inadequately segmented networks.

A tale of two teams

As organizations increasingly adopt cloud strategies, their internal IT and security teams have evolved. For many organizations, the evolutionary process can be quite challenging and painful. Different teams often have different priorities and internal conflict ensues.

Cloud teams generally consist of cloud infrastructure architects and engineers, cloud security architects and engineers, and DevOps/DevSecOps. Some key challenges for the cloud team include the following:

- » **Pressure to meet deadlines and maintain compliance:** DevOps/DevSecOps teams are under constant pressure to meet business goals in an agile fashion, while staying compliant with government regulations (from organizations such as the Department of Health and Human Services [HHS], the Food and Drug Administration [FDA], and the Securities and Exchange Commission [SEC]), as well as internal and industry compliance requirements.
- » **Complexity of the cloud:** The cloud is complex, whether the DevOps/DevSecOps team is migrating existing applications or developing new applications. In addition, the organization may be using multiple clouds, and every cloud is different, from network design to compute, storage, network, and shared security responsibility architectures.
- » **Lack of security expertise:** The DevOps/DevSecOps team is not always aware of advanced threats that can cause a breach, and the security products that can prevent those advanced threats. Frequently, the DevOps/DevSecOps team will lean on native cloud service provider (CSP) security for default “good enough” that are ineffective against today’s modern cyber threats.



WARNING

For the cloud team, networking — and, by extension, network security — is not always part of their worldview. This can result in bypassing learning about, interacting with, and owning these products. Thus, cloud teams typically prefer to allow another team, or the CSP, to address network security needs.

Traditional network security teams typically consist of network security architects and engineers, as well as a head of infrastructure. Some common challenges encountered by these teams in cloud and virtualized environments include the following:

- » **Loss of control:** The network security team is losing control as applications migrate to the cloud or new applications are built in the cloud. The cloud team has more authority, and they come up with the blueprint for migrating or building applications for the cloud. Often, security is an afterthought in these application blueprints.

- » **Lack of visibility:** Network security teams lack visibility into threats in multi-cloud environments. New network connections may be untracked and not compliant with internal security guidelines. Hybrid cloud and multi-cloud environments can lead to multiple point security solutions that create security gaps, as well as gaps in infrastructure visibility, due to the lack of a single pane of glass to easily correlate threat information from different sources.
- » **Fighting the perception that network security is too hard, and not “worth it”:** When the network security team recommends a network security solution, they’re left with the burden of proof to show that their recommendation won’t slow down the business and that it’s worth the additional effort as compared to the easy and “good enough” native CSP security.
- » **Lack of cloud expertise:** Network security teams frequently lack cloud infrastructure and/or cloud application expertise. They often don’t have the skills to evaluate network security products on the basis of integration with automated network security provisioning tools to keep pace with the application development life cycle.



REMEMBER

The network security team is still held accountable for breaches but doesn’t always have the authority, control, and visibility to address security issues in the cloud let alone across multiple clouds. These issues are often exacerbated by pushback from cloud teams concerned about missing deadlines.

Security of the cloud versus security in the cloud

In a traditional on-premises data center or private cloud model, you’re responsible for the security of the entire stack from the physical hardware and connectivity to the applications and data — and everything in between. You’re also responsible for the physical security (and resilience) of the data center itself, including the building, electricity, cooling, and internet connectivity.

Things are a bit different in the public cloud. The shared responsibility model is a relatively straightforward framework for

defining which elements in the technology stack for a particular public cloud service — that is, IaaS, PaaS, or SaaS — are the responsibility of the CSP and which elements are the responsibility of the customer. Although the framework itself is fairly easy to understand, the actual implementation of the shared responsibility model — including the different solutions available in the cloud to secure your services — can be quite confusing and creates security gaps even within a single cloud environment.

For example, most CSPs provide some level of data resiliency, such as disk mirroring to prevent data loss in the event of a disk failure, VM high availability prevents data loss in the event of a hard failure of a server, and availability zones to prevent data loss in the event of a data center failure. However, these data protection measures are implemented to ensure the CSP can meet their service-level agreements (SLAs) and to protect the CSP. You're still responsible for backing up your data. If your data is encrypted due to a ransomware attack, your CSP will not pull a copy of your data from a different availability zone for you. Instead, you must restore your data from a good backup that you — not the CSP — created.

Across all service models (IaaS, PaaS, and SaaS), the CSP is responsible for the management and security of the platform foundation, as well as the networking, storage, compute, and virtualization elements of the technology stack. In an IaaS model, the customer is responsible for managing and securing everything else — operating systems, middleware, runtime, applications, and data. At the other end of the service model spectrum, SaaS customers are only responsible for their data, and the CSP takes care of everything else (see Figure 1-4). The CSP secures the platform foundation. This encompasses the hardware and software that provide the networking, storage, computing, and virtualization services — as well as standard operating systems, such as Red Hat Enterprise Linux (RHEL) and Windows Server.



TIP

Another way to think about cloud security and the shared responsibility model is that the CSP is responsible for security *of* the cloud, while the customer is responsible for security *in* the cloud (see Figure 1-5).

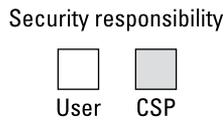
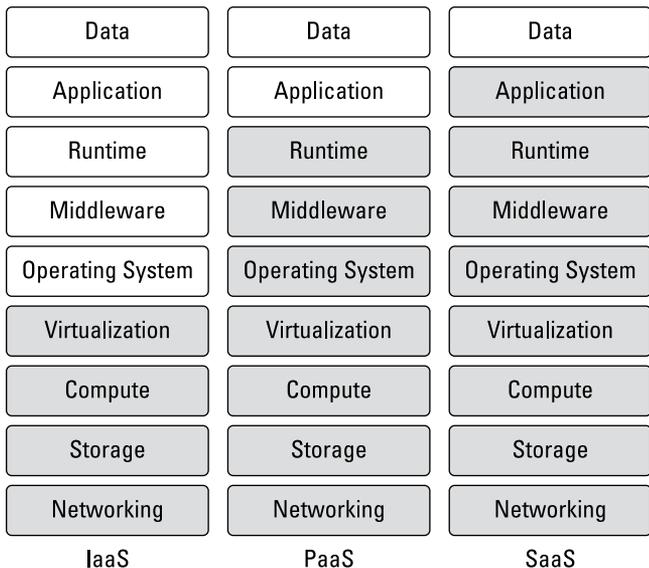


FIGURE 1-4: The shared responsibility model.

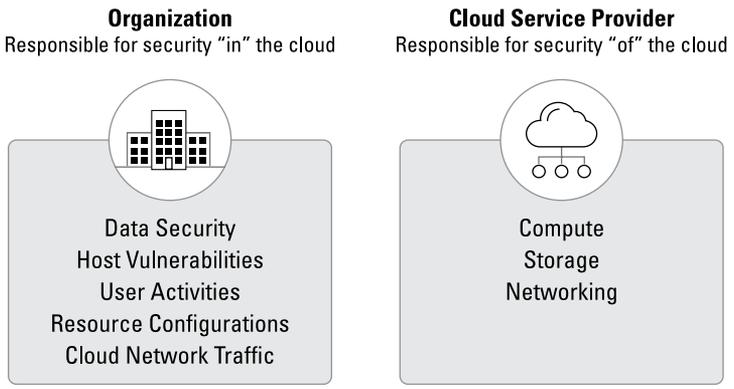


FIGURE 1-5: Division of security responsibilities in the public cloud.

Compliance requirements

The regulatory landscape continues to grow ever more complex with security and privacy regulations — such as the U.S. Health Insurance Portability and Accountability Act (HIPAA), Payment Card Industry Data Security Standards (PCI DSS), European Union (EU) General Data Protection Regulation (GDPR), and Brazilian General Data Protection Law (LGPD) — being regularly passed and amended in virtually every country, state, and jurisdiction around the world.

Moving applications and data from an on-premises data center to the public cloud can significantly impact an organization's compliance strategies. On the one hand, data residency and single tenancy requirements, among others, can complicate compliance strategies. However, in many cases, you can benefit from the security controls the CSP uses on its own infrastructure and the certifications your CSP has attained. However, this benefit is not “carte blanche.” *Remember:* In the shared responsibility model (discussed earlier in this chapter), you're always responsible for the security (and privacy) of your data.



REMEMBER

An effective strategy for compliance in cloud and virtualized environments requires organizations to centralize security management so they can align policies across the entire environment, including on-premises, branch, remote, virtualized, private and public cloud, hybrid cloud, multi-cloud, and edge environments.

To address modern network security challenges, organizations need to adopt a Zero Trust strategy (see Figure 1-6) that extends across the entire enterprise attack surface including on-premises and cloud environments. I explain how to get started with Zero Trust in Chapter 2.

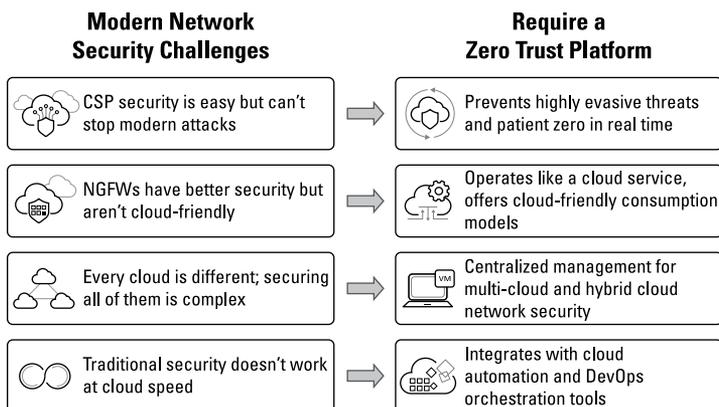


FIGURE 1-6: Modern network security challenges require a Zero Trust platform.

IN THIS CHAPTER

- » Defining Zero Trust
- » Getting started with a Zero Trust architecture and software firewalls
- » Protecting cloud applications with software firewalls

Chapter 2

Understanding Software Firewalls and Zero Trust

In this chapter, you learn about Zero Trust and least-privilege access, how software firewalls help you achieve a Zero Trust strategy, and how to extend Zero Trust to virtualized and cloud (public, private, hybrid) environments with software firewalls.

What Is Zero Trust?

Zero Trust is a strategic approach to cybersecurity that eliminates implicit trust, continuously validates user and object identities, and enforces least-privilege access throughout a digital session. Based on the principle of “never trust, always verify,” Zero Trust is designed to protect modern environments and enable digital transformation by using strong authentication methods, leveraging network segmentation, preventing lateral movement, providing Application Layer (Layer 7) threat prevention, and creating simple, intuitive, yet powerful granular, least-privilege access policies.



REMEMBER

In a Zero Trust architecture, every device, user, application, workload, and network flow is considered to be a potential threat and, therefore, must be continuously authenticated and authorized, and digital transactions must be inspected during the session.

WHAT IS LEAST-PRIVILEGE ACCESS?

A security best practice when configuring roles and permissions for any software environment is to apply least-privilege access. In other words, when you adhere to the principle of least privilege, you focus on ensuring that no user or group has access rights or permissions that exceed the minimum required to perform their role within the organization.

The benefits of least privilege

The main benefit of least privilege is that it limits the potential damage caused by a security breach. In an environment where a user has access to more resources than they need, anyone who manages to compromise that user account will likewise have access to those systems. But by restricting access to the minimum permissions necessary, you limit the impact of a breach.

Least-privilege access can also simplify audits. When you follow the principle of least privilege, you can perform audits of your access policies to determine whether any policies give users more access rights than they require. You can then take steps to address the risk.

An example of least-privilege access

To understand what least privilege means in practice, consider a cloud environment that is shared by multiple users within an organization. Some of the users are developers, while others are IT engineers. The developers use one set of development/test virtual machines (VMs) to build and test applications. The IT engineers use another set to deploy applications for production use.

To configure least-privilege access in this scenario, you would configure cloud identity and access management (IAM) roles and policies in such a way that the developers could create, modify, and run only the specific VMs they use for development/test purposes. Likewise, the IT engineers would only be able to access production VMs.

The opposite of least privilege in this example would be to create IAM rules that give all team members access to all VMs. You can assume that developers may sometimes need to access production VMs, and

IT engineers may sometimes want to see what's happening in the development/test environment. However, this approach would increase the potential impact of a security breach. If a developer's account is compromised, for instance, the attackers would be able to access all VMs in the environment if the account has access to all of them. With least-privilege access in place, only the dev/test environments would be exposed.

Zero Trust principles address the outdated assumption that everything inside an organization's network should be implicitly trusted, upon which traditional perimeter-based security models operate. This implicit trust means that when they're on the network, users — including threat actors and malicious insiders — are free to move laterally and access or exfiltrate sensitive data due to a lack of granular security controls.

Many organizations today spend a lot of time and effort attempting to identify, and reduce, their constantly expanding attack surface. This challenge has become exponentially more difficult with the proliferation of the Internet of Things (IoT), mobile devices, and remote working.

Instead of trying to manage a constantly growing attack surface, Zero Trust principles focus on the protect surface, which consists of the organization's most critical resources, defined by one or more of the following elements:

- » Users
- » Applications
- » Infrastructure

An organization's protect surface is orders of magnitude smaller than its attack surface and is far easier to determine (see Figure 2-1). Instead of creating a single perimeter at a macro level to protect an entire attack surface, Zero Trust establishes multiple micro-perimeters as close as possible to the protect surface. These micro-perimeters move with the protect surface through segmentation gateways that restrict access to known allowed traffic and applications.

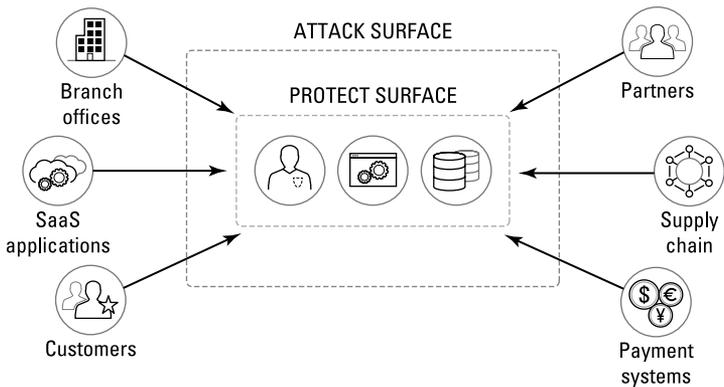


FIGURE 2-1: Zero Trust principles shrink the number of potential entry points for cyberattacks.



TIP

With digital transformation accelerating in the form of a growing hybrid workforce, continued migration to the cloud, and the transformation of security operations, taking a Zero Trust approach has never been more critical. If done correctly, a Zero Trust architecture results not only in higher overall levels of security, but also in reduced security complexity and operational overhead.

How Do Software Firewalls Help You Achieve Zero Trust?

A Zero Trust strategy requires micro-perimeters to be established as close as possible to users, applications, and infrastructure, creating a well-defined and manageable protect surface to enforce security. Of course, the protect surface can be highly ephemeral and mobile, particularly when it comes to modern, cloud applications leveraging microservices, containers, and elastically scalable architectures. Traditional, perimeter-based physical firewall appliances are, thus, ill-suited for a robust Zero Trust strategy.

Software-based next-generation firewalls (NGFWs) can be deployed as part of a Zero Trust strategy to establish micro-perimeters around dynamic protect surfaces across an enterprise's different environments. This includes physical and virtualized data centers; branch and remote locations; and private clouds, public clouds, hybrid clouds, edge clouds, and multi-clouds.



TIP

Until recently, NGFWs were deployed solely as physical appliances. However, you can't deploy hardware appliances in virtualized and public cloud environments or move them dynamically with your critical data, assets, applications, and services. Software firewalls are the ideal solution for a Zero Trust strategy for all clouds — including public cloud, private cloud, hybrid cloud (on-premises and cloud), multi-cloud environments.

Extending Zero Trust Practice to Cloud Applications

Traditional, network firewalls were built on physical hardware appliances and designed to provide perimeter-based security between the “trusted” corporate network and the “untrusted” internet. Much has changed over the past two decades, and this notion of implicit trust and well-defined perimeters is no longer valid — in fact, it's even dangerous. Additionally, traditional physical hardware firewalls can't be deployed in a virtualized or cloud environment.

As previously discussed, Zero Trust practices require an organization to identify discrete protect surfaces and implement micro-segmentation between their users, applications, and infrastructure. Even within an environment limited to a small on-premises data center, deploying physical firewalls to protect every user, application, and infrastructure component would be impractical. In the public cloud, it's impossible — you can't deploy your own hardware in the cloud.

Software firewalls offer a practical solution that enables organizations to implement a Zero Trust strategy in their on-premises data centers and extend Zero Trust to their cloud-based applications as well.



REMEMBER

Software firewalls address the issue of not being able to deploy hardware in the public cloud, but also address the ephemeral and transient nature of cloud and cloud-native applications (and their associated microservices) which may exist for only a few seconds or minutes, and can dynamically move to different VMs, containers, or regions — and even different clouds.

IN THIS CHAPTER

- » Getting started with virtual firewalls
- » Leveraging cloud firewalls
- » Extending software firewalls to Kubernetes environments with container firewalls
- » Enhancing your capabilities with cloud-delivered security services

Chapter 3

Discovering Software Firewall Types

In this chapter you learn about the different types of software firewalls — including virtual, container, and cloud firewalls — that are available to support a wide array of unique business requirements and use cases.

Virtual Firewalls

Virtual next-generation firewalls (NGFWs) provide all the capabilities of physical hardware firewalls in a virtual machine (VM) form factor. Some key capabilities and features of virtual firewalls include the following:

- » **Consistently address cloud security requirements.** Protect operating systems, platforms, access, control, data, and more from known and unknown threats, to meet your security obligations in a cloud service provider's shared responsibility model (see Chapter 1), and maintain internal, industry, and government compliance.

- » **Secure virtualized resources and hypervisors.** Block lateral cyber-threat movement between applications and workloads with segmentation, micro-segmentation, and continuous traffic inspections in virtualized and cloud environments.
- » **Simplify management.** Isolate and protect critical systems with consistent threat prevention and security policies delivered from the same console across all your cloud environments including private cloud, public cloud, multi-cloud, and hybrid cloud.
- » **Scale security at the speed of business.** Integrate automated security provisioning directly into DevOps workflows and continuous integration/continuous deployment (CI/CD) pipelines to support business agility.

US SIGNAL SAFEGUARDS SENSITIVE INFORMATION

US Signal is a leading provider of data center and cloud services. With eight data centers in the Midwest, the company hosts cloud solutions, provides colocation space, and delivers best-of-breed security services powered by its own secure, robust fiber network.

Challenge

In 2020, US Signal decided to expand its cloud footprint and data protection capabilities. The company had seen 300 percent growth in its data center business over the previous five years and sought to meet increasing demand for its services from customers during a period when many businesses were migrating to the cloud to support remote work.

At the time, US Signal was operating multiple firewall platforms from several vendors. As part of its expansion, US Signal wanted to consolidate platforms and work with a single vendor. Doing so would mean its engineers wouldn't need to learn the ins and outs of multiple systems and would gain a centralized perspective of the company's own security infrastructure.

The company also wanted to take advantage of automation to further curtail the likelihood of human error, reduce overhead, and lift some of the burden on its engineers. Provisioning firewalls for customers was a

cumbersome, time-consuming process. By leveraging automation, the team would be able to do more with less and expedite deployments.

Requirements

US Signal evaluated all major security vendors. To even be considered, a solution had to operate in a VMware ESXi for vSphere bare-metal hypervisor environment. Best-in-class security was a primary requirement because the chosen solution would protect both US Signal's infrastructure and its customers'. Scalability was equally critical in order to keep up with the company's appetite for expansion. In addition, US Signal judged real-time updates as essential to equip the virtual firewalls with the latest security features and threat intelligence. Ease of automation was of tremendous importance, because the company deploys virtual firewalls for hundreds of customers. Licensing and price point were also key.

Solution

US Signal chose Palo Alto Networks VM-Series NGFWs to provide complete Zero Trust network security for both its internal IT infrastructure and the product offerings that US Signal deploys to customers. By taking advantage of Palo Alto Networks automated provisioning and integration with Ansible and Terraform orchestration products, US Signal deployed every virtual firewall with a full suite of cloud-delivered security services. These services included GlobalProtect to safeguard mobile users, WildFire advanced malware analysis, threat prevention, advanced Uniform Resource Locator (URL) filtering, and Domain Name System (DNS) security to secure traffic and prevent sophisticated known and unknown attacks.

Impact

Deploying Palo Alto Networks VM-Series NGFWs made a significant impact with benefits that included the following:

- Firewall provisioning time cut by 97 percent through extensive automation
- Customer purchase decisions expedited with feature-rich virtual NGFWs
- Customer trust enhanced with advanced cloud security
- Protection against zero-day threats with uninterrupted operations
- Competitive advantage and continued growth with unsurpassed security solutions

Cloud Firewalls

Cloud firewalls are delivered with many of the same features of a physical or virtual NGFW but deployed and configured in minutes with just a few familiar cloud-native developer clicks. By leveraging firewalls as cloud services, organizations can benefit from cost savings by eliminating the need to install or maintain security hardware or software firewalls across their entire organization.

Cloud firewalls are virtual firewall-powered and provide best-in-class security for network security teams and deeply integrated into CSP environments for cloud-native ease of use for DevOps teams (see Figure 3-1).

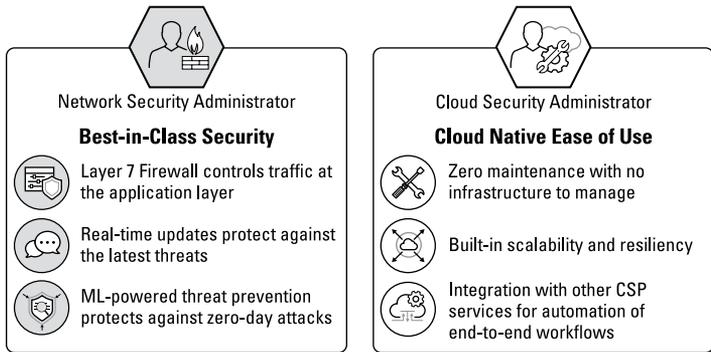


FIGURE 3-1: Modern enterprises need both best-in-class security and cloud-native ease of use.

The cloud firewall approach enables organizations to:

- » Gain complete visibility into and control over their networks without having to deploy physical appliances, thereby reducing support costs
- » Automatically deliver scalability and resilience with no infrastructure to manage
- » Take advantage of a cloud-native ease-of-use service with built-in real-time high efficacy and scale against modern cyberattacks

For example, Palo Alto Networks provides these cloud firewalls in partnership with CSPs, integrated software vendor (ISV) partnerships called Cloud NGFWs, for use in AWS and Microsoft Azure environments, and white-labeled OEM partnerships with Google Cloud and Oracle Cloud Infrastructure.

So, how do these cloud firewalls differ from other software firewalls? For one thing, they are managed services procured directly in the marketplaces where CSPs list and provide offerings from ISVs that run as native services, which make them easy to deploy, operate, and automate within those environments. Not only that, but once procurement decisions have been made, users then remain in those portals to quickly set up and deploy the cloud firewalls, along with items such as rulestacks and automated security profiles. Cloud-native user interfaces can complete the deployment and configuration process in just a few minutes.

When these cloud firewalls are tightly integrated with the CSP, users gain access to security features purpose-built for those particular environments. As you'll read in the next section about cloud-delivered security services, those capabilities can include:

- » Helping to stop zero-day threats in real time
- » Securing applications as they connect to legitimate web-based services with inline deep-learning security
- » Controlling traffic and allowing only sanctioned applications to traverse the network with Layer 7 (the application layer) classification
- » Stopping vulnerability exploits and sophisticated attacks, unknown exploits, spyware, as well as malware and command-and-control (C2) communications
- » Preventing and detecting file-based threats by using dynamic, static, and machine learning analysis techniques
- » Countering the most advanced DNS-layer attacks from exploiting customer networks and stealing their data

Once deployed within the CSP environment, cloud firewalls address the very real security needs of where organizations deploy applications and store critical data in the public cloud. These software firewalls, of course, bolster security and integrate with built-in CSP services for automation and scale. What's more, cloud firewalls enable simple and consistent firewall policy management across the locations where applications reside.

Because companies often have applications on-premises as well as in public clouds, advanced cloud firewalls can come with management systems that provide comprehensive network security across these environments. This ensures organizations don't need separate consoles for separate locations. Many organizations have spent years building up their security posture on-premises and understandably resist the idea of starting over with new and isolated policies and processes in the cloud. When done right, unified management can centrally manage all of their NGFWs across all environments, including particular CSPs.

Container Firewalls

Containerization and orchestration technologies are used extensively in the cloud to deliver modern cloud-native applications at scale. Traditional network firewalls were not designed to handle the requirements of these extremely large, complex, and agile environments.



Kubernetes is a popular open-source container orchestration system. In Kubernetes, containers run on resource nodes, which can be either physical or, more commonly, VMs. Namespaces are ways to isolate resources in Kubernetes. They can isolate pods, as well as services, storage, and annotations. Developers rarely have to deal with nodes explicitly, but nodes impact how firewalls operate. Firewalls running outside the managed Kubernetes cluster can't determine which container pod is the source of outbound traffic because all source IP addresses are translated to the node IP address. Thus, to a firewall outside the Kubernetes cluster, all outbound traffic from the node looks the same. This is somewhat analogous to doing network address translation (NAT) on a network switch, thus converting all private Internet Protocol (IP) addresses to a single public IP address, before the traffic hits the firewall (see Figure 3-2).

Additionally, east-west traffic between microservices within a Kubernetes environment never leaves the cluster perimeter; as such, the firewall doesn't inspect traffic between pods.

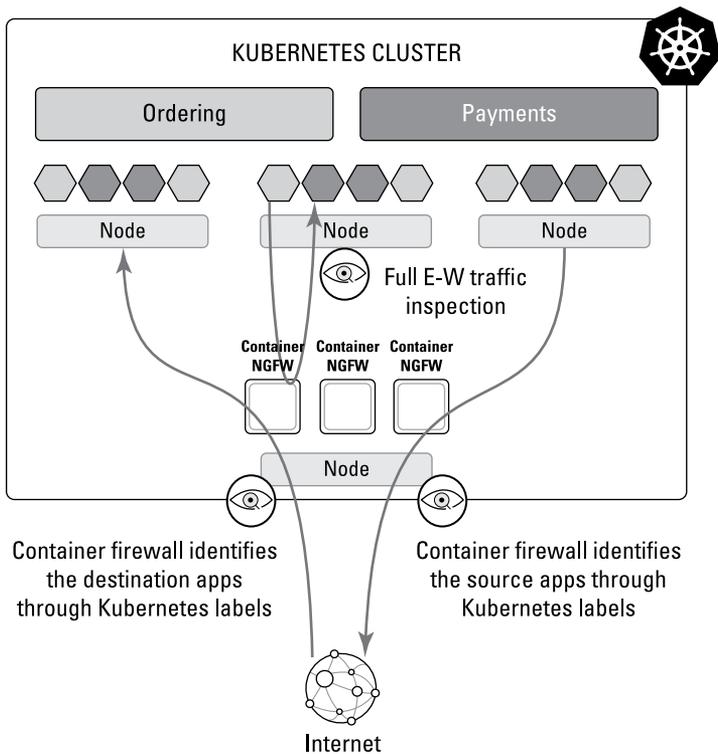


FIGURE 3-2: Due to the use of NAT in Kubernetes, all outbound traffic is identified by the node source IP address.

Kubernetes creates challenges for traditional security tools, but it also presents new opportunities to enhance security. For example, Kubernetes namespaces help simplify cluster management by making it easier to apply certain policies to specific parts of a cluster. Security teams can use namespaces to isolate workloads and reduce the risk of an attack spreading within a cluster.

Namespaces can also be used to establish resource quotas to mitigate potential damage from a cluster breach impacting other containers and to limit malware taking advantage of misconfigured containers.

A complete security solution for containerized environments requires the ability to secure traffic that crosses namespace boundaries in a Kubernetes environment, as well as traffic that travels outbound to other cloud and on-premises workloads, such as VMs, bare-metal servers, and external code repositories like GitHub. Additionally, Kubernetes pods can also serve as applications for clients (inbound traffic). Virtual firewalls outside the cluster, for example, won't be able to apply granular traffic to each application/namespace uniquely. This capability requires information about the internal state of objects such as namespaces, pods, and containers. However, this information isn't available outside the Kubernetes environment. Thus, to meet the Zero Trust architecture principle of protecting the workloads as close to the load as possible, the only effective solution is to deploy a container firewall solution *within* Kubernetes environments (see Figure 3-3).

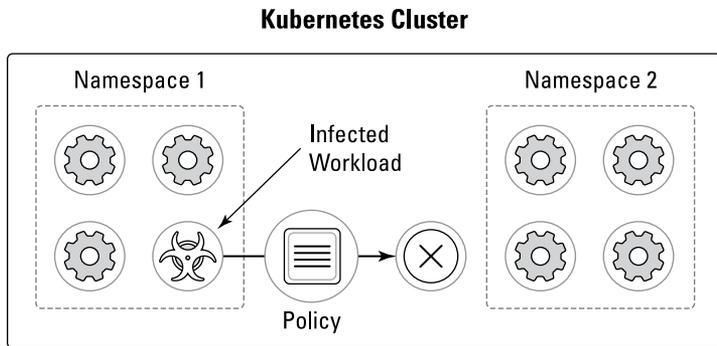


FIGURE 3-3: Security policies based on namespaces prevent the spread of an attack within a cluster.



TIP

Clustered deployments are best suited for large Kubernetes environments where a distributed deployment would be resource intensive and cost prohibitive.



REMEMBER

Native integration with Kubernetes enables container firewalls to leverage contextual information about the containers in the environment to create security policies. For instance, container namespaces can be used to define a traffic source in a firewall policy.

Cloud-Delivered Security Services

For cybersecurity to be everywhere, it has to be delivered as a cloud-based service. When this happens, cybersecurity solutions can take advantage of the elasticity of the cloud to provide scalability that supports today's rapidly changing businesses, as well as globally crowd-sourced threat intelligence and machine learning (ML) to protect customers from known and unknown threats in real time.



TIP

Palo Alto Networks cloud-delivered security services are natively integrated, offering consistent best-in-class artificial intelligence (AI)- and ML-enabled protection everywhere. Backed by the Palo Alto Networks Unit 42 threat research team and a network of more than 85,000 global customers, intelligence is shared from all threat vectors to stop known, unknown, and zero-day threats 180 times faster than other solutions. Key cloud-delivered security services from Palo Alto Networks include the following:

- » **Advanced threat prevention service** prevents all known threats and zero-day exploits across all traffic in a single pass.
- » **Advanced WildFire sandboxing** deals with unknown threats.
- » **Advanced URL filtering service** enables safe access to the internet for users in any location by blocking access to known and new malicious websites.
- » **DNS security service** disrupts attacks that use DNS for command-and-control (C2) traffic and data theft.
- » **Internet of Things (IoT) security service** provides visibility, prevention, and enforcement for IoT devices and operational technology (OT).
- » **Software as a service (SaaS) security service** delivers complete visibility and granular enforcement across all user, folder, and file activity for SaaS applications.
- » **Enterprise data loss prevention (DLP) service** discovers, monitors, and protects all sensitive data and support compliance across every network, cloud, and user.
- » **Software-defined wide-area network (SD-WAN) service** simplifies management, enables granular application-defined policies, and implements secure cloud-enabled branches.

IN THIS CHAPTER

- » Securing public cloud infrastructure
- » Protecting your private cloud
- » Deploying software firewalls in hybrid cloud and multi-cloud environments
- » Extending software firewalls to branch locations
- » Defending 5G networks

Chapter 4

Exploring Software Firewall Use Cases

In this chapter, you explore common software firewall use cases including public cloud, private cloud, hybrid cloud, and multi-cloud environments; virtualized branches; and 5G networks.

Public Cloud

Public cloud service providers (CSPs) sell resources such as computing cycles and storage bytes on a pay-as-you-go basis. The hardware and software that make up the platform are invisible to you — your team only sees a black-box abstraction that supports your applications and data. The benefits are substantial: You avoid capital investment and costly refresh cycles and gain virtually unlimited scalability, high reliability, dependable data backup, and business agility.

The following sections explore some of the key security use cases and challenges in public cloud environments and how software-based next-generation firewalls (NGFWs) — including virtual firewalls, container firewalls, and cloud firewalls — solve these challenges.

Application security detects hard-to-find threats

Port-based Internet Protocol (IP) firewalls and network security groups implemented by CSPs lack application-level visibility into network traffic and have limited threat prevention capabilities. As a result, native cloud security groups won't discover threats that exploit open ports (for example, 80 and 443) or target vulnerabilities in non-web apps.

Software firewalls deployed in the public cloud as virtual firewalls, container firewalls, or cloud firewalls, inspect every inbound packet and block suspicious traffic based on application type, content type, or user/device identity, going beyond simple port blocking to protect traffic over open ports. These software firewalls also provide advanced security capabilities, such as intrusion prevention system (IPS) and sandboxing, to defend against both known and unknown vulnerabilities at the edge of and within a public cloud environment.



REMEMBER

Software NGFWs have all the capabilities of physical NGFWs. Plus, they can automatically follow the dynamic nature of virtualized and containerized applications and workloads within public cloud, private cloud, and hybrid cloud environments as application demand scales up and down.

Outbound traffic protection stops exfiltration

To exfiltrate data from a target environment after it has been breached, attackers often take advantage of encrypted traffic flows that are permitted by the organization, such as Secure Sockets Layer (SSL) or Transport Layer Security (TLS) encryption, to secure data as it leaves the environment. For example, after gaining access to your environment, an attacker identifies valuable information and uses a Domain Name System (DNS) tunneling technique to exfiltrate it from the compromised application by hiding the data in encrypted DNS/C2 traffic.



TIP

Software firewalls can decrypt traffic for outbound content inspection. The DNS security service (discussed in Chapter 3) ensures that even allowed encrypted traffic flows are inspected and protected.

Filtering and inspection boosts developer security

Native CSP firewall offerings typically have limited capabilities to filter and inspect outbound traffic leaving the cloud environment. As a result, if developers download compromised open-source code from a public code repository, they may unwittingly introduce malware into the development environment. After these vulnerabilities are introduced inside application code, attackers can move laterally to locate information for exfiltration.



TIP

Cloud-delivered advanced Uniform Resource Locator (AURL) filtering services (discussed in Chapter 3), deployed on software firewalls in the public cloud, ensure that developers can only access known good repositories that are maintained and secured internally (see Figure 4-1).

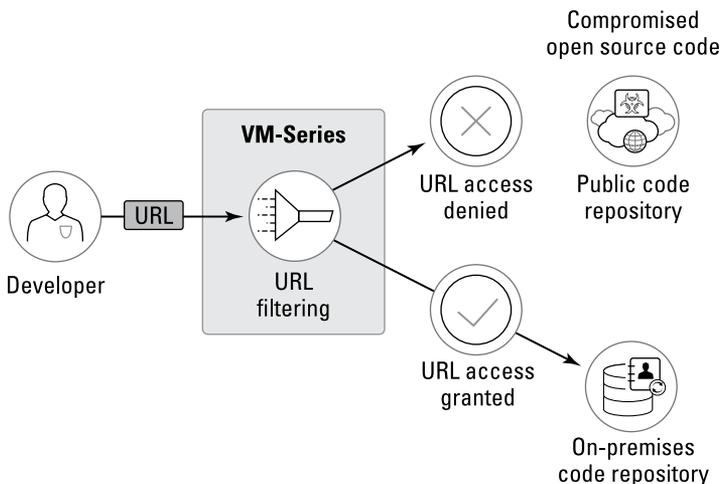


FIGURE 4-1: Advanced URL filtering services prevent developers from accessing compromised code in a public repository.

Private Cloud

For organizations that require more control of their applications and data than the public cloud provides, a private cloud delivers many of the same benefits without multitenancy, data residency, and other issues that exist in the public cloud. This is why many companies host their sensitive data and core business applications in private clouds and use public clouds to provide additional scalability.



REMEMBER

Some of the most critical applications to the enterprise are core business applications, such as the following:

- » **Enterprise resource planning (ERP) and manufacturing resource planning (MRP)** where critical business planning and financials reside
- » **Product line management (PLM)** where business-critical intellectual property and product design details are captured and stored
- » **Quality management systems (QMSs)** where business and product quality issues are tracked that businesses certainly wouldn't want exposed to competitors
- » **Business intelligence (BI)/analytics** (for example, Tableau) and any associated databases the BI tools connect to
- » **Human resources management systems (HRMSs)** that contain employee personally identifiable information (PII), organizational structure, compensation, and so on

Compliance considerations for security regulations and standards — such as the Payment Card Industry Data Security Standard (PCI DSS), the Health Insurance Portability and Accountability Act (HIPAA), generally accepted accounting principles (GAAP), the Sarbanes–Oxley Act (SOX), and so on — take top security priority in a private cloud environment to protect the core business.

Unique core applications and workloads that are generally more predictable in usage are also potential candidates for virtual data centers and private cloud environments.



TIP

According to the *Flexera 2022 State of the Cloud Report*, 84 percent of enterprises use at least one private cloud.

In the following sections, we explore some of the key security use cases and challenges in private cloud environments and how software firewalls solve these challenges.

Segmentation and micro-segmentation protect against lateral movement

Network security teams often lack visibility and control over east–west traffic (that is, inter–virtual machine [VM] and intra–VM traffic that doesn’t cross a perimeter firewall) in virtualized networks and private clouds, leading to relatively flat and open virtual infrastructure where any workload can communicate with any other workload. Attackers take advantage of this fact to move laterally within a breached network, establishing persistence, escalating privileges, and looking for additional targets with little risk of detection (see Figure 4–2).

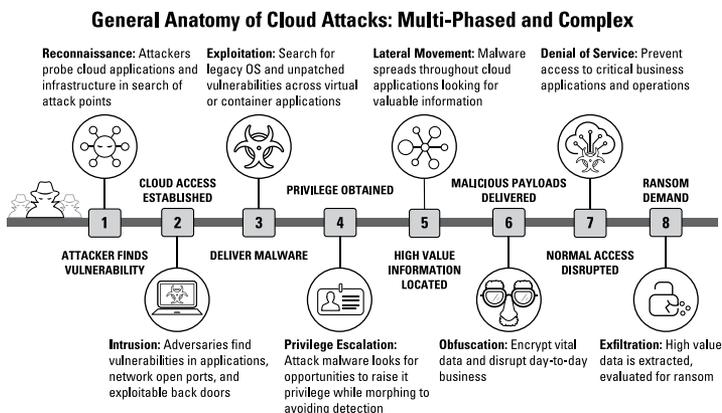


FIGURE 4-2: The general anatomy of cloud attacks.



TIP

Software firewalls, such as virtual firewalls, deployed strategically throughout a virtualized, private cloud environment create trust zones (or, as discussed in the previous chapter, smaller Zero Trust protection surfaces) based on an organization’s risk profile and tolerance level. Critical applications are placed in their own trust zones, with threat prevention services turned on to inspect traffic to and from the application. Sensitive data subject to compliance standards is enclosed in another trust zone that enforces the security measures required for compliance. These same security controls can be deployed in containerized environments, using container firewalls deployed natively within a Kubernetes environment.

Augment software-defined networking with threat prevention

Many organizations deploy software-defined networking (SDN) solutions — examples include VMware NSX, Cisco Application-Centric Infrastructure (ACI), and Nutanix Flow — to implement segmentation in their virtualized environments. SDN solutions are good at restricting traffic based on predefined policies, but they lack the ability to provide modern security enforcement and automatically detect threats in allowed traffic flows.

For example, an organization might deploy an SDN solution to simplify networking in a virtualized environment and implement micro-segmentation for some of their critical applications. SDN restricts workload communication to only the traffic necessary for the applications to function. However, in addition to accessing shared databases, the applications leverage application programming interfaces (APIs) and other shared interfaces as part of their normal operations, communications, and transactions. An attacker can use these allowed connections to move laterally or exfiltrate information from the environment. Software firewalls can seamlessly integrate with VMware NSX, Cisco ACI, Nutanix Flow, and other SDN solutions to allow the rapid and accurate insertion of advanced security services, such as intrusion prevention or sandboxing, between micro-segments to inspect and protect allowed traffic.

VDI security addresses threats to remote and distributed workforces

Virtual desktop infrastructure (VDI) deployments offer a host of operational efficiencies, but also present challenges for network security teams that lack adequate visibility and control. Threats entering the network via compromised VDI devices can move laterally and target other high-value private cloud assets.

For example, an organization deploys VDI in their data center or private cloud. Because end users control these internet-connected machines, the VDI end points have a low trust level and must be segmented from the rest of the environment. Any allowed traffic to shared services or applications must be inspected for threats.



TIP

Virtual NGFWs deployed at the edge of the VDI environment ensure that virtual desktops are properly segmented and traffic is inspected appropriately (see Figure 4-3).

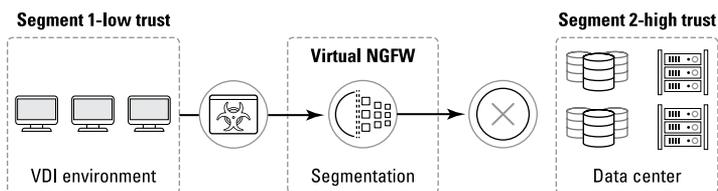


FIGURE 4-3: Segmentation and threat prevention capabilities to protect data center and private cloud assets from attacks originating in the VDI environment.

Hybrid Cloud and Multi-Cloud

According to the *Flexera 2022 State of the Cloud Report*, multi-cloud is now the de facto standard, with 89 percent of enterprises having a multi-cloud strategy and 80 percent having a hybrid cloud strategy. As organizations divide application hosting between multiple public clouds and private clouds (45 percent of applications are siloed on different clouds according to the Flexera report), overall security posture becomes more fragmented and difficult to manage. Each part of the environment requires its own policy model and security controls, which increases operational complexity, creates security gaps, and causes delays for cloud migration initiatives.

For example, a large enterprise might run its mission-critical, cloud-native applications in a private cloud and two different public cloud environments. To enforce consistent security policies across all three parts of this hybrid environment, the security team must duplicate policies across three clouds using the native controls in each — a labor-intensive and error-prone task, made all the more challenging by the highly dynamic and ephemeral nature of cloud and virtual environments. Managing overall security posture requires the team to develop expertise in each cloud’s controls and management interface, as well as Kubernetes for container applications — no small feat.

This challenge becomes even more difficult in hybrid environments that extend to on-premises data centers, potentially requiring network security teams to manage hardware and software firewalls across multiple public and private clouds, as well as on-premises data centers and branch offices.



TIP

Network security NGFWs — including the software firewalls (virtual, container, and some cloud firewalls) that are deployed in multi-cloud and hybrid environments — can all still be managed from the same console. This lets security teams deliver the same best-in-class security capabilities to each environment and extend a uniform policy model across the entire ecosystem to ensure consistency and complete visibility that simplifies your overall security posture.

Virtualized Branch

Many organizations are embracing digital transformation for their branch offices, retail locations and even remote critical infrastructure, effectively creating software-defined branches. Software firewalls are ideal to help these organizations enforce segmentation in their branches and use secure software-defined wide-area networking (SD-WAN) for branch connectivity.

The following sections explore some of the key security use cases and challenges in virtualized branch locations and how software firewalls can help solve these challenges.

Achieving compliance with local branch segmentation

Regulatory compliance mandates often require segmentation between sensitive applications and data. For example, the Payment Card Industry Data Security Standard (PCI DSS) requires segmentation of the cardholder data environment (CDE) to protect sensitive payment card information. Added challenges for branch locations often include a lack of local IT and security resources and limited space for deploying physical hardware.



TIP

Software NGFWs can be deployed on existing servers in branch locations, typically called universal customer premise equipment (uCPE), avoiding the need for additional hardware. This allows security teams to create a single security policy to enforce required segmentation and intrusion prevention to protect the security of data as required by any applicable regulatory compliance requirements. The organization can centrally manage its distributed branch security, using the same management console that controls security for its data center, private cloud, and public cloud networks (see Figure 4-4).

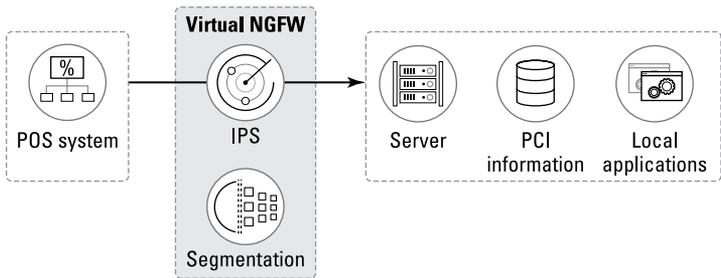


FIGURE 4-4: Support for compliance frameworks (for example, PCI DSS) with IPS, segmentation, and other security features.

Software-based perimeter security simplifies deployment

With the growing adoption of software as a service (SaaS) and other remote services, many organizations are turning to SD-WAN to simplify networking and reduce the hardware needed at branch locations. Traditional hardware firewalls require on-site expertise to deploy and configure, but branches often lack local technical expertise and extra space for hardware.

Software NGFWs can replace hardware firewalls in branch and remote locations while delivering the same level of security as at the corporate headquarters and consolidating services with less hardware. Virtual NGFWs can be deployed as VMs on existing servers or uCPE, alleviating the need to ship, install, and maintain separate hardware firewalls. This saves the costs of truck rolls and adding space-consuming appliances. Virtual NGFWs can also deliver SD-WAN to further consolidate branch connectivity and security (see Figure 4-5).

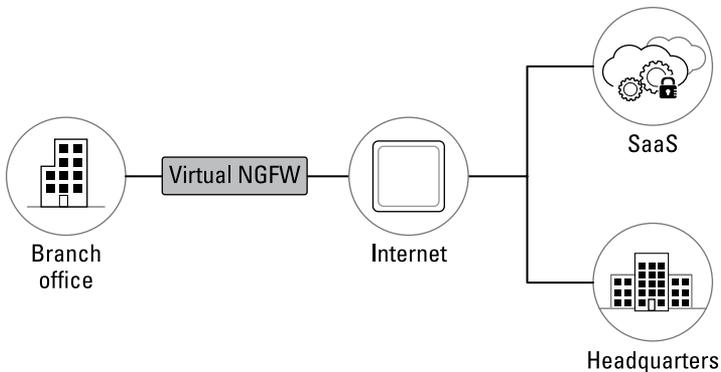


FIGURE 4-5: Perimeter deployment to protect the branch from internet-based threats.

Secure SD-WAN increases performance and network return on investment

As more applications are hosted in clouds, traditional wide-area networks (WANs) are becoming suitable to effectively manage networks and security. Leveraging legacy connectivity options like Multiprotocol Label Switching (MPLS) to backhaul branch traffic to a central data center is not efficient in today's cloud-connected world.

Virtual NGFWs can be used to build a hub-and-spoke SD-WAN deployment between branch locations, the core data center, public cloud environments, and SaaS applications. SD-WAN can also be deployed as a full mesh with both hub-to-spoke and branch-to-branch connectivity (see Figure 4-6).

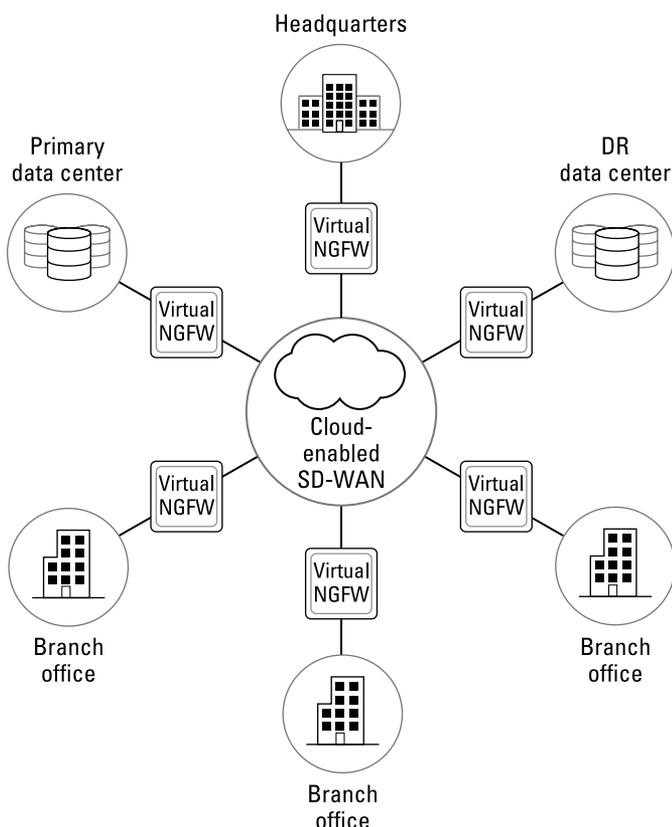


FIGURE 4-6: Virtual NGFWs in a hub-and-spoke architecture.

5G creates disruptive business opportunities for service providers and enterprises because it can move beyond delivering connectivity and enable them to use security as a business enabler and competitive advantage. The evolution of 5G opens the door to exciting new services, but it also increases the number of potential intrusion points, amplifying the security impact. To tap into the 5G business opportunities with minimal risk of being exploited by threat actors, organizations need end-to-end visibility and automated security across all network locations.



TIP

Palo Alto Networks 5G-native security helps organizations safeguard their networks, customers, and clouds, and extend Zero Trust to their 5G environments. 5G-native security supports all key mobile infrastructure environments — including on-premises, virtual, and containerized — across public and private telco clouds, and multi-access edge computing (MEC) environments with machine learning (ML)-powered NGFWs.



WARNING

The fundamental shift in 5G network architectures further intensifies the impact on the security landscape, with growth in the number of intrusion points, including attacks inside mobile tunnels, and threats within apps traversing cellular traffic. This larger attack surface increases the need for application-aware Layer 7 security to detect known and unknown threats across all network locations and all signaling traffic.

IN THIS CHAPTER

- » Stopping zero-day threats and implementing least privilege
- » Delivering a consolidated, future-proof security platform
- » Maximizing flexibility and simplifying management
- » Securing any network and all clouds
- » Working with your existing tools and accelerating security posture
- » Delivering a high return on investment

Chapter 5

Ten Questions to Ask Your Software Firewall Vendor

Here are ten important questions to help you evaluate potential software firewall vendor solutions for your organization.

Do They Stop Zero-Day Threats?

Waiting to protect against zero-day threats only after the first victim has been compromised can lead to rapid lateral spread, putting your whole organization at risk. Some solutions try to prevent zero-day attacks by holding files for analysis, but this ultimately doesn't work because it harms the user experience and delays business.



TIP

Look for a software firewall that protects against zero-day threats in real time with built-in artificial intelligence (AI) and deep machine learning (ML) and is constantly refreshed with frontline intelligence from cloud-delivered security services.

Do They Provide Least-Privilege Access Control?

Least-privilege access control ensures that users, applications, and/or devices have only the minimum permissions necessary to perform an authorized task or action. Least privilege is critical to implementing a Zero Trust strategy. For example, you may only allow a shipping and receiving clerk with a smart handheld scanner running a receiving application to access the enterprise resource planning (ERP) system for shipping receipt transactions. Or, you may restrict (block access for) all mobile devices into the corporate human resources management system (HRMS) and only allow HR and executives to access it through their on-premises desktop computers.



TIP

Ensure your software firewall provides continuous trust verification and security inspection to extend Zero Trust to all cloud applications and workloads. The combination of application, user, device, and content identification capabilities in a next-generation firewall (NGFW), along with intuitive security policies, provides a fantastic level of least-privilege access control.

Can Security Be Consolidated into a Security Platform?

A security platform approach consolidates all critical security functions — including centralized management; AI-driven, ML-powered NGFW; a security operating system with ML and analytics; and cloud-delivered security services.



TIP

A single-pass processing architecture is critical in an NGFW to deliver high throughput and low latency.



TIP

Consolidating security services can save you time in securing infrastructure and deliver a higher return on investment (ROI) when consolidating security functions. Make sure that security features like intrusion prevention systems (IPSS), sandboxing, Uniform Resource Locator (URL) filtering, Domain Name System (DNS) security, data loss prevention (DLP), Internet of Things (IoT) security, and software-defined wide-area networking (SD-WAN) capabilities are fully integrated into the platform.

Can They Provide Consistent and Future-Proof Protection?

Consuming cybersecurity innovation can be difficult. Organizations waste time deploying additional hardware or software every time they want to take advantage of a new security technology. They invest more resources managing their ever-expanding security infrastructure instead of improving their security controls to stay ahead of attackers and prevent threats.



TIP

Your software firewall should enable teams to quickly discover, evaluate, and use new security technologies. Security teams should be able to collaborate between different apps, share threat context and intelligence, drive automated response and enforcement with deeply integrated applications, and provide consistent protection of application traffic across all clouds — including hybrid cloud, multi-cloud, and virtualized environments. This way, they can solve the most challenging security use cases with the best technology available — today and in the future — and they can do so without the cost or operational burden of deploying new infrastructure.

Can You Dial Up/Down Security with Flexible Consumption?

A key benefit of the cloud is flexible, pay-as-you-go subscription-based consumption plans. Your software firewall should offer this same business flexibility so you can rapidly scale security needs up/down as requirements evolve and change.

Do They Provide Single-Pane-of-Glass Management?

Look for a software firewall that provides a single pane of glass from which to manage all your firewalls regardless of form factors and locations from on-premises to clouds. This reduces complexity by simplifying the configuration, deployment, and management of your security policies. You want a tool that correlates firewall logs to provide application, network and security insights, as well as surface malicious behavior, which can often get buried in the noise or lost in the security gaps of point products missed across shared responsibility security boundaries.

Can They Secure Any Cloud and Application Architecture Model?

Data and applications reside everywhere — in your network and in the cloud. According to the *Flexera 2022 State of the Cloud Report*, 89 percent of enterprises use multiple public, private, and/or hybrid clouds — more than five different clouds on average. Compounded with software as a service (SaaS) environments, organizations must now secure sensitive data in the network and a variety of cloud environments. In addition, legacy security tools and techniques designed for static networks don't work with cloud-native tools or capabilities. Moreover, native security services from the cloud providers themselves, such as Amazon Web Services (AWS), Microsoft Azure, and Google Cloud Platform (GCP), typically provide only Layer 3/4 protections and are specific to that cloud provider.

To succeed, your organization needs cloud security that extends policy consistently from the network to the cloud, stops malware from accessing and moving laterally (east–west) within the cloud, simplifies management, minimizes the security policy lag as virtual workloads change, and closes security gaps that can result in hybrid cloud/multi-cloud environments with point products. Your software firewall must protect the resident applications and data

with the same security posture that you may have established on your physical network. To secure multi-cloud deployments, the firewall must support a variety of cloud and virtualization environments, including all major public cloud providers and virtualized private clouds. The firewall must integrate with native cloud services, such as Amazon and Azure, and automation tools, such as Ansible and Terraform, to fully integrate security with the way your application developers work in your cloud-first application architecture model.



REMEMBER

With Palo Alto Networks, you have the ability to truly protect all clouds and any virtualized environment with deep integrations that accelerate your deployments and get you to a consistent security posture faster. Palo Alto Networks provides deep integrations, including the following:

» **Cloud service providers**

- Alibaba Cloud
- AWS
- GCP
- IBM Cloud
- Microsoft Azure
- Oracle Cloud

» **Kubernetes containers**

- Amazon Elastic Kubernetes Service (EKS)
- Azure Kubernetes Service (AKS)
- Google Kubernetes Engine
- OpenShift
- Rancher
- VMware Tanzu

» **Software-defined networks and hypervisors**

- Cisco Application Centric Infrastructure (ACI)
- Linux Kernel-based Virtual Machine (KVM)
- Microsoft Hyper-V

- Nutanix Acropolis Hypervisor (AHV)
- Nutanix Flow
- OpenStack
- VMware ESXi
- VMware NSX

Do They Work with Your Automation/Orchestration Tools?

Your software firewalls deployed in the cloud must work the way your application developers work. This means integrating with the automation and orchestration tools that these teams use and are familiar with to rapidly deliver new software releases.



TIP

Palo Alto Networks software firewalls support automated deployment, scaling, and policy changes with native capabilities, including dynamic address groups (DAGs) and application tags, as well as popular automation and orchestration tools: Ansible, AWS CloudFormation, Azure Resource Manager (ARM) templates, Helm Charts, Kubernetes, Terraform, and more.

Are They Proven to Accelerate Security Posture?

Software firewalls must be simple to deploy yet powerful enough to address the unique security challenges in hybrid cloud and multi-cloud environments to reduce downtime and increase staff productivity.



TIP

According to the Forrester September 2021 report, *The Total Economic Impact of Palo Alto Networks VM-Series Virtual Firewalls*, deploying VM-Series virtual firewalls resulted in an average 30 percent reduction in time required to achieve proper security posture and savings of \$436,760 over three years.

Do They Have a Track Record of Delivering High ROI?



TIP

Palo Alto Networks VM-Series virtual firewalls pay for themselves and can provide savings measured in millions of dollars over a short period of time. According to the Forrester September 2021 report, *The Total Economic Impact of Palo Alto Networks VM-Series Virtual Firewalls*, they resulted in the following benefits:

- » **115 percent** ROI over three years with a six-month payback period.
- » **90 percent** reduction in time required to deploy firewalls.
- » **80 percent** improved network and security team efficiencies, potentially saving \$1.3 million over three years.
- » **67 percent** decrease in employees experiencing a downtime event — only 6 percent are impacted by these incidents.

Glossary

AI: *See* artificial intelligence (AI).

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

Application Layer: Layer 7 of the OSI model, responsible for identifying and establishing availability of communication partners, determining resource availability, and synchronizing communication. *See also* Open Systems Interconnection (OSI) model.

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.

attack surface: The sum total of the devices and connections that threat actors could potentially use to penetrate network defenses.

botnet: A broad network of malware-infected end points (bots) working together and controlled by an attacker through C2 infrastructure. *See also* command-and-control (C2).

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

CI/CD: *See* continuous integration/continuous delivery (CI/CD).

cloud firewall: Joint cloud service provider and security solution provider with cloud-native user-interface ease of use powered by virtual next-generation firewall (NGFW) for best-in-class security against known and unknown cyber threats. The joint solution may be managed as an independent software vendor (ISV) or original equipment manufacturer (OEM) and managed by the cloud service provider (CSP) solution. Palo Alto Networks ISV solutions include Cloud NGFW for AWS and Cloud NGFW for Azure, powered by VM-Series virtual firewalls. Palo Alto

Networks OEM solutions include Google Cloud IDS and Oracle OCI Network Firewall, powered by VM-Series virtual firewalls. *See also* cloud service provider (CSP), container firewall, independent software vendor (ISV), original equipment manufacturer (OEM), software firewall, *and* virtual firewall.

cloud service provider (CSP): A third-party company offering cloud-based platform, infrastructure, application, and/or data storage 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.

container firewall: A category of software firewalls that provide next-generation firewall (NGFW) capabilities in an application container form factor deployed and managed within Kubernetes environments and all clouds — private, public, and hybrid — that contain container applications managed with Kubernetes or cloud adaptations of Kubernetes (for example, Amazon EKS, Azure Kubernetes Services [AKS], Google Kubernetes Engine, OpenShift, Rancher, VMware Tanzu, and so on). The Palo Alto Networks container firewall solution is CN-Series Container Next-Generation Firewall. *See also* cloud firewall, software firewall, *and* virtual firewall.

continuous integration/continuous delivery (CI/CD): A DevOps environment supported by automation such that changes to application source code and infrastructure configuration are built, integrated, and deployed automatically. *See also* DevOps.

CSP: *See* cloud service provider (CSP).

DAG: *See* dynamic address group (DAG).

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

DDoS: *See* distributed denial-of-service (DDoS).

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

DevOps: The culture and practice of improved collaboration between software developers and IT operations.

distributed denial-of-service (DDoS): An attack in which the attacker initiates simultaneous denial-of-service attacks from many systems (potentially tens of thousands), typically bots in a botnet, with the intention of making the system or network unavailable for use. *See also* botnet.

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).

dynamic address group (DAG): A dynamic address group populates its members dynamically using lookups for tags and tag-based filters. Dynamic address groups are very useful if you have an extensive virtual infrastructure where changes in virtual machine (VM) location or IP address are frequent. *See also* virtual machine (VM) *and* Internet Protocol (IP).

enterprise resource planning (ERP): Business management software used to collect, store, manage, and interpret data from many business activities.

ERP: See enterprise resource planning (ERP).

Federal Information Security Modernization Act (FISMA): A U.S. federal law that defines a comprehensive framework to protect government information, operations, and assets.

FISMA: See Federal Information Security Modernization Act (FISMA).

GDPR: See General Data Protection Regulation (GDPR).

General Data Protection Regulation (GDPR): A data privacy law that strengthens data protection requirements for European Union (EU) residents and addresses the export of personal data outside the EU.

Health Insurance Portability and Accountability Act (HIPAA): A U.S. federal act that addresses security and privacy requirements for medical systems and information.

HIPAA: See Health Insurance Portability and Accountability Act (HIPAA).

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.

hypervisor: In a virtualized environment, the supervisory program that controls allocation of resources and access to communications and peripheral devices.

laaS: See infrastructure as a service (IaaS).

IAM: See identity and access management (IAM).

identity and access management (IAM): A software service or framework that allows organizations to define user or group identities within software environments and then associate permissions with them.

independent software vendor (ISV): A software producer that is not owned or controlled by a hardware manufacturer (or in cases outlined in this book, the cloud service provider); a company whose primary function is to distribute software.

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

Internet of Things (IoT): The network of physical, connected objects embedded in electronics, operating systems, software, sensors, and network connectivity.

Internet Protocol (IP): The OSI Layer 3 (Network) 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.

IoT: See Internet of Things (IoT).

IP: See Internet Protocol (IP).

IPS: See intrusion prevention system (IPS).

ISV: See independent software vendor (ISV).

least-privilege access: Access in which a user or object is assigned only the minimum level of permissions needed to perform an authorized task.

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).

MNO: See mobile network operator (MNO).

mobile network operator (MNO): A provider of wireless communications services that owns or controls all the elements necessary to sell and deliver services to an end user. Also known as a wireless service provider, wireless carrier, cellular company, or mobile network carrier.

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 necessarily 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 you know) and a one-time passcode sent to a mobile phone that has previously been registered with the organization (something you 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.

NAT: See network address translation (NAT).

network address translation (NAT): The process of converting internal, privately used IP addresses in a network to external, public IP addresses. See *also* Internet Protocol (IP).

Network Layer: Layer 3 of the OSI model, responsible for routing and related functions that enable data to be transported between systems on the same network or on interconnected networks. See *also* Open Systems Interconnection (OSI) model.

next-generation firewall (NGFW): A network security platform that fully integrates traditional firewall and network IPS capabilities with other advanced security functions that provide DPI for complete visibility, accurate application, content and user identification, and granular policy-based control. See *also* deep packet inspection (DPI) *and* intrusion prevention system (IPS).

NGFW: See next-generation firewall (NGFW).

OEM: See original equipment manufacturer (OEM).

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

original equipment manufacturer (OEM): Traditionally defined as a company whose goods are used as components in the products of another company, which then sells the finished item to users. The second firm is referred to as a value-added reseller (VAR) because, by augmenting or incorporating features or services, it adds value to the original item. The VAR (cloud service providers) works closely with the OEM (Palo Alto Networks), which often customizes designs based on the VAR company's needs and specifications.

OSI: See Open Systems Interconnection (OSI) model.

PaaS: See platform as a service (PaaS).

Payment Card Industry Data Security Standard (PCI DSS): A standard set of requirements developed for the protection of personal data related to credit, debit, and cash card transactions.

PCI DSS: See Payment Card Industry Data Security Standard (PCI DSS).

personally identifiable information (PII): Information (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.

PII: See personally identifiable information (PII).

platform as a service (PaaS): A cloud-based service model 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 are 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.

protect surface: The data, applications, assets, services, and infrastructure that must be protected in a Zero Trust architecture. See also Zero Trust.

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

radio access network (RAN): Part of a mobile telecommunication system that connects cellular wireless-capable devices (such as a mobile phone) to a public and/or private mobile core network via an existing network backbone.

RAN: *See* radio access network (RAN).

SaaS: *See* software as a service (SaaS).

SDN: *See* software-defined network (SDN).

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

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 network (SDN): An approach to networking that separates the network control and management processes from the underlying hardware and makes them available as software.

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

software firewall: A family of next-generation firewalls (NGFWs), including virtual firewalls, container firewalls, and cloud firewalls, specifically designed to be deployed with DevOps agility into virtual environments and all clouds — private, public, and hybrid clouds. *See also* cloud firewall, container firewall, DevOps, *and* virtual firewall.

Transport Layer: Layer 4 of the OSI model, responsible for providing transport and end-to-end transmission control. *See also* Open Systems Interconnection (OSI) model.

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).

VDI: *See* virtual desktop infrastructure (VDI).

virtual desktop infrastructure (VDI): A desktop operating system running within a virtual machine (VM) on a physical host server. *See also* virtual machine (VM).

virtual firewall: A category of software firewalls that provide next-generation firewall (NGFW) capabilities in a virtual machine (VM) form factor that can be deployed into virtual environments and all clouds — private, public, and hybrid. May also be the baseline NGFW for cloud firewalls that are deeply integrated with cloud-native user interfaces for ease of use and deployments. The Palo Alto Networks virtual firewall solution is the VM-Series Virtual Next-Generation Firewall. *See also* cloud firewall, container firewall, *and* software firewall.

virtual machine (VM): An instantiation of an operating system running within a hypervisor. *See also* hypervisor.

VM: *See* virtual machine (VM).

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.

Zero Trust: A strategic approach to cybersecurity that eliminates implicit trust, continuously validates user and object identities, and enforces least-privilege access. *See also* least-privilege access.

Take a platform security approach in the cloud

Modern businesses need security solutions that are simple to use and don't slow down cloud development teams. These same solutions need to ensure network security teams can confidently deploy and manage them to secure the business from increasingly sophisticated threats. Software firewalls deliver the same robust capabilities as hardware-based next-generation firewalls (NGFWs) and can be deployed in on-premises, branch, public cloud, private cloud, and hybrid cloud/multicloud environments. *Software Firewalls For Dummies* shows you how to secure any network and all clouds.

Inside...

- Execute a Zero Trust strategy
- Enforce least-privilege access
- Protect containerized environments
- Simplify multicloud security
- Protect remote branch locations
- Secure cloud applications



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 coauthor 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: 978-1-394-16555-1
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.