

LEARNING MADE EASY

Palo Alto Networks Special Edition

Data Center & Hybrid Cloud Security

for
dummies[®]
A Wiley Brand



Address
security challenges

—
Implement
best practices

—
Enact threat
prevention

Brought to
you by



Lawrence C. Miller

About Palo Alto Networks

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



Data Center & Hybrid Cloud Security

Palo Alto Networks Special Edition

by Lawrence C. Miller

for
dummies[®]
A Wiley Brand

Data Center & Hybrid Cloud Security For Dummies®, Palo Alto Networks Special Edition

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

Copyright © 2020 by John Wiley & Sons, Inc., Hoboken, New Jersey

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

Trademarks: Wiley, For Dummies, the Dummies Man logo, Dummies.com, 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. Palo Alto and the Palo Alto logo are trademarks or registered trademarks of Palo Alto Networks, Inc. All other trademarks are the property of their respective owners. John Wiley & Sons, Inc., is not associated with any product or vendor mentioned in this book.

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

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

ISBN 978-1-119-68155-7 (pbk); ISBN 978-1-119-68156-4 (ebk)

Manufactured in the United States of America

10 9 8 7 6 5 4 3 2 1

Publisher's Acknowledgments

We're proud of this book and of the people who worked on it. For details on how to create a custom *For Dummies* book for your business or organization, contact info@dummies.biz or visit www.wiley.com/go/custompub. For details on licensing the *For Dummies* brand for products or services, contact BrandedRights&Licenses@Wiley.com.

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

Project Editor: Martin V. Minner

Associate Publisher: Katie Mohr

Editorial Manager: Rev Mengle

Business Development

Representative: Karen Hattan

Production Editor: Umar Saleem

Table of Contents

INTRODUCTION	1
About This Book	2
Icons Used in This Book.....	2
CHAPTER 1: The Evolution of the Data Center	3
Understanding the Impact of the Cloud on the Data Center.....	3
Recognizing Security Challenges in the Data Center and Hybrid Cloud	6
Looking at Why Legacy Security Infrastructure is Ineffective	8
Firewalls	8
Intrusion prevention.....	9
Proxies.....	9
Defining Security Requirements.....	10
Gain complete visibility	10
Minimize the attack surface.....	11
Automate threat protection	11
CHAPTER 2: Security Challenges in Hybrid Clouds	15
Cloud Security and The Shared Responsibility Model	15
The Dynamic Nature of Modern Threats	21
Ransomware.....	24
Credential theft	25
DNS-based attacks.....	27
Targeted “Low-and-Slow” Attacks and APTs.....	31
CHAPTER 3: Delivering Consistent Security Using Zero Trust	33
Gaining Complete Visibility	34
Application identification	35
User identification.....	37
Content identification.....	38
Minimizing Your Attack Surface with Segmentation.....	39
Using Dynamic Security to Support Moves and Changes	41

CHAPTER 4: Leveraging Unmatched Threat Protection 45

- The Challenges of Defense in Depth..... 45
- What Is a Perimeter?..... 47
- Threat Protection Components..... 48
 - Endpoint protection..... 48
 - Security orchestration and automation..... 52
 - Cloud-based threat intelligence..... 56

CHAPTER 5: Ten Evaluation Criteria for Network Security in the Data Center and Hybrid Cloud Environment 57

- Safe Enablement of Applications in the Hybrid Cloud..... 58
- Identify Users and Enable Appropriate Access..... 60
- Comprehensive Threat Protection..... 61
- Flexible, Adaptive Integration 63
- Secure Access for Mobile and Remote Users 63
- One Comprehensive Policy, One Management Platform 64
- Cloud Ready 65
- Automate Routine Tasks and Focus on the Threats That Matter 66
- Flexible Deployment Options 67
- Consume New Innovations Easily in a Broad Partner Ecosystem 68

GLOSSARY 69

Introduction

For decades, traditional data centers have been the most predictable and controlled infrastructure to house high-risk and proprietary assets, such as personal information, medical records, and financial information.

Today, data is everywhere. Modern application workloads are moving across multiple data centers as well as private, public, and hybrid clouds, anywhere around the globe, based on business requirements. These highly distributed applications need to dynamically grow, shrink, and move, and be rewritten and redeployed according to business needs.

The data center is also evolving to allow fast and flexible application deployment. To support this level of application elasticity and mobility, enterprises are transforming their data centers with a modern architecture. A modern data center utilizes technologies such as virtualization, cloud, and software-defined networking to deliver application workloads everywhere across physical data centers as well as hybrid and multi-cloud environments.

A modern infrastructure allows your organization to extend data centers into cloud services. This evolution of the data center enables flexible scaling for network, storage, and compute demand surges. A modern data center offers the best of both worlds: security, performance, and reliability with agility, scalability, and cost savings across on-premises data centers and multiple public, private, and hybrid clouds.

The modern data center helps IT organizations deliver greater business opportunities but also introduces new risks. Data centers that span multi-cloud environments offer a larger attack surface, which can translate to increased complexity in networking and cybersecurity. It is critical to maintain full visibility and precise control of your data center regardless of the architecture. It is also important to implement a best practice methodology and adopt an approach to data center security that is independent of the individual environments you're utilizing.

About This Book

Data Center & Hybrid Cloud Security For Dummies helps you rethink your approach to security to better protect your data and workload in the on-premises data center and across multi-cloud environments. The chapters in this book explore:

- » The evolution of the traditional data center to a modern data center architecture (Chapter 1)
- » Security challenges in the hybrid cloud (Chapter 2)
- » Security best practices for the data center and hybrid cloud (Chapter 3)
- » The role of threat prevention in the data center and hybrid cloud (Chapter 4)
- » Key evaluation criteria for network security in the data center and hybrid cloud (Chapter 5)

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

Icons Used in This Book

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



REMEMBER

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



TECHNICAL
STUFF

If you seek to attain the seventh level of NERD-vana, perk up! This icon explains the jargon beneath the jargon and is the stuff nerds are made of.



TIP

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



WARNING

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

IN THIS CHAPTER

- » Looking at enterprise cloud adoption
- » Understanding hybrid cloud and data center security challenges
- » Recognizing the limitations of legacy security infrastructure
- » Taking a new approach to security in the hybrid cloud

Chapter 1

The Evolution of the Data Center

This chapter explores how public and private cloud adoption have led to a hybrid cloud environment in the enterprise, the unique security challenges and limitations of legacy security, and how to address security in the data center and hybrid cloud.

Understanding the Impact of the Cloud on the Data Center

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

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

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

Virtualization enables organizations to utilize data center infrastructure more effectively, which helps lower costs and improves operational efficiencies through automation, agility, and reusability of compute and network resources. Virtualization initiatives often begin with server consolidation projects to more effectively utilize server hardware resources. These efforts often expand beyond host virtualization to include virtualization of storage, networking, and other physical infrastructure, in order to realize other virtualization benefits. In this way, enterprises are increasingly able to run their data center operations more efficiently — effectively transforming their data centers.

Few organizations today can afford to ignore public cloud offerings, and rarely do an organization’s physical data centers go away altogether, because it’s neither feasible nor desirable to adopt a cloud strategy based solely on the public cloud. Instead, many organizations are adopting a hybrid cloud model to leverage the advantages of both public and private multi-cloud computing.

The main driver for moving to a hybrid cloud strategy is business operations. Such a strategy enables IT organizations to better support constantly and rapidly changing business conditions and new opportunities, by being more flexible and agile. Other business benefits include the following:

» **Supports legacy applications:** Some organizational applications may be legacy in nature and therefore do not lend themselves to either a lift-and-shift or refactoring of the workloads for use in the public cloud, thereby permanently relegating them to the on-premises data center.

- » **Improves performance:** Many applications experience unpredictable demand. Virtualization technologies, such as resource schedulers and virtual machine (VM) migrations, provide intelligent, automated placement of application workloads. This prevents resource and network-throughput contention issues and maximizes server utilization by moving virtual workloads to underutilized resources within the data center.
- » **Supports mobility:** The network link between the data center and the back office is critical, but it's not the only connection. Increasingly, an organization's customers and users rely on high-speed Internet access to reach the data center from all sorts of devices including the Internet of Things (IoT), industrial control systems (ICS), and 5G. Moving applications to the cloud enables organizations to better support their customers, as well as remote and mobile users, by placing applications and their associated data "closer" to those customers and users.
- » **Reduces time-to-market:** Dynamic provisioning of on-demand resources in the cloud reduces time-to-market by enabling new applications to be delivered more rapidly.
- » **Promotes standardization:** Virtualization technologies enable organizations to create standard server builds and easily clone standard configurations.
- » **Increases scalability:** Compute, storage, and networking resources can quickly and easily be scaled up or down to support changing business requirements, such as mergers and acquisitions, as well as cyclical business environments.



REMEMBER

The modern, hybrid cloud enables greater IT efficiency, automation, and agility, supporting the delivery of new application workloads across dynamic network fabric and virtual machine infrastructure. A hybrid cloud strategy allows organizations to create the right mix of cloud and traditional IT to suit their needs.

Recognizing Security Challenges in the Data Center and Hybrid Cloud

The data center is rapidly evolving from a traditional, closed environment with static, hardware-based computing resources to one in which there is a mix of traditional and cloud computing technologies (see Figure 1-1).

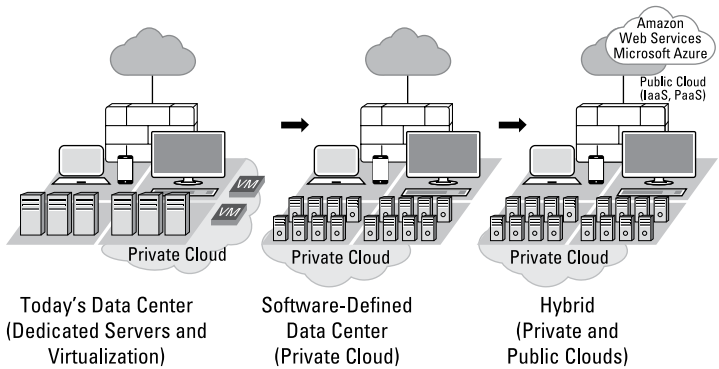


FIGURE 1-1: Data centers are evolving to include a mix of hardware and cloud computing technologies.

Three distinct trends and challenges are emerging in the modern data center (see Figure 1-2):

» **New data center technologies:** Big data analytics and virtualized applications are producing an explosion of data within the data center. To accommodate this transformation, enterprises are adopting technologies including virtualization, software-defined networking (SDN), and hyperconverged infrastructures, that allow for the extension of new and existing applications into hybrid and multi-cloud deployments. The elastic nature of these workloads requires the network to scale up and out on demand, but this makes it difficult for data center teams to enforce security as workloads move across data centers and clouds. Data center teams can't see where the workloads are, who is using them, or where those users are connecting from. These teams are also challenged to enforce proper access controls and implement a Zero Trust architecture. Fundamentally, security must be consistently enforced regardless of workload dynamics.

» **Changes in application development:** Changes in DevOps and adoption of continuous integration/continuous deployment (CI/CD) pipelines are enabling business agility and growth. Software release schedules have moved from once or twice a year to many times per day. New application technologies, such as containers, microservices, and Representational State Transfer (REST) application programming interfaces (APIs), are changing the way applications are designed and deployed. Moreover, these new applications no longer reside in a single data center. Instead, they have components that run on-premises, in the cloud, across hybrid clouds, and even on end-users' browsers. This new application development approach places great pressure on network teams to ensure their networks are robust, matching the agility and speed of DevOps as well as the line of business.

» **Orchestration of workloads that are everywhere:** As organizations race to the cloud to gain an economic advantage, data centers and clouds are often indistinguishable. Workloads can be orchestrated and provisioned across the cloud, and applications can reside on-premises one day and in the cloud the next day. This workload mobility is an underlying driver of the hybrid cloud architecture. New orchestration tools allow workloads to be easily provisioned, deployed, and administered across multiple data center and cloud environments. Network operations for the hybrid cloud require full, continuous network insight to better manage security across the data center footprint as well as support orchestration of workloads across multiple locations and in a multi-cloud infrastructure.

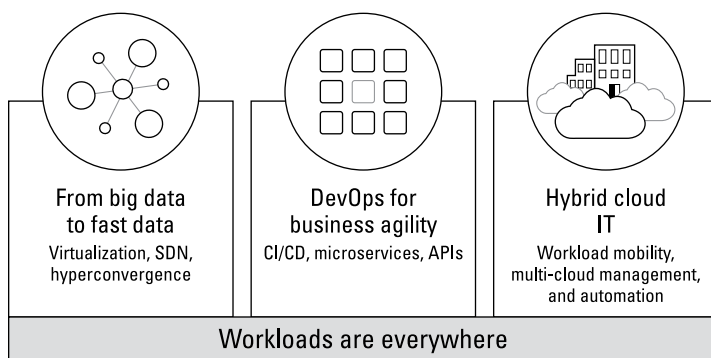


FIGURE 1-2: Dynamic changes in the data center.

Looking at Why Legacy Security Infrastructure is Ineffective

Legacy security infrastructures are generally flat network architectures that rely on a perimeter firewall as their only point of traffic inspection and control. These traditional port-based firewalls provide limited value in a cloud and mobile world where network boundaries have all but disappeared and the majority of traffic in a data center is east-west (traffic flow within the data center). For security to be effective, it must deliver perimeter security as well as build trust zones within an organization's internal network. This ensures that traffic between apps and services of different trust levels is filtered using best-in-breed network security services, such as intrusion prevention and Domain Name System (DNS) security. The same level of protection must extend to public clouds to ensure consistent network security and segmentation in hybrid environments as well.

In addition to limited visibility into network traffic context, many of these solutions apply static policies and controls based on more or less permanent physical and/or logical locations, such as IP addresses and ports. They are thus unable to adapt effectively to hybrid cloud environments in which application workloads have shorter lifecycles and can move dynamically between on-premises, private, and public cloud locations.



TIP

A next-generation firewall must deliver the same capabilities consistently at the network edge or to aid segmentation in the cloud. It is important to move security controls as close as possible to the workloads they are protecting. Different physical and virtual form factors allow organizations to accomplish this goal across a hybrid cloud model.

Firewalls

Firewalls are often used as a first line of defense, but legacy port-based firewalls provide only coarse filtering of traffic and limited network segmentation. One drawback to port-based firewalls is that they use protocol and port to identify and control what gets in and out of the network. This port-centric design is ineffective when faced with malware and evasive applications that hop from port to port until they find an open connection to the network. Such firewalls themselves have little ability to identify and control advanced threats.

Within a hybrid cloud environment, the majority of all network traffic today consists of east-west communications between servers in the data center, many of which are virtual machines. Thus, legacy port-based firewalls are largely ineffective because the traffic never passes through the firewall. Attackers are free to move laterally throughout the data center with little risk of detection. Modern security must be deployed strategically to address both public and private cloud attack vectors, in order to provide comprehensive protection of the organization's systems and data in a hybrid cloud environment. This requires strategic placement of next-generation firewalls throughout both on-premises and private cloud environments, as well as within public cloud environments to filter and inspect all inbound, outbound, and east-west traffic.



REMEMBER

Security policies should be based on the identity of users and the applications in use — not just on IP addresses, ports, and protocols. Without knowing and controlling exactly who (users) and what (applications and content) has access to the network and its various assets, data centers and hybrid cloud environments may be compromised by threats that can easily bypass port-based network controls.

Intrusion prevention

Traditional IPS solutions use a mix of exploit-based signatures — which can be produced quickly but provide limited coverage — and vulnerability-based signatures — which take longer to create but provide coverage for a broad range of exploits — and attempt to apply the appropriate signatures to specific types of traffic, based on port. This limitation means that malware or exploits on unexpected or nonstandard ports are likely to be missed. Additionally, IPS solutions lack the depth of exploit detection needed to protect hybrid clouds — most IPS solutions only look for a few hundred types of common exploits — well short of the tens of thousands that exist.

Proxies

Proxy solutions are another means of network traffic control. But they too look at a limited set of applications or protocols and only see a partial set of the network traffic that needs to be monitored. By design, proxies need to mimic the applications they are trying to control so they struggle with updates to existing applications

and new applications. As a result, although proxies understand a few protocols in depth, they typically lack the breadth of protocol support needed to control the tunnels and protocols within protocols that attackers use to hide their true traffic.

Defining Security Requirements

Data center network teams protecting hybrid clouds with traditional security approaches face a security complexity trifecta:

- » Limited visibility and imprecise control
- » An ever-expanding attack surface
- » Increasingly advanced cyberthreats

Overcoming these challenges requires adopting a new hybrid cloud protection methodology to gain complete visibility, minimize the attack surface, and automate threat protection (see Figure 1-3).

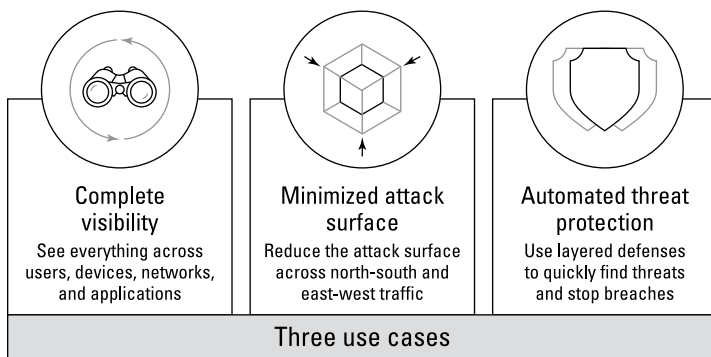


FIGURE 1-3: What's needed for hybrid cloud security.

Gain complete visibility

Complete visibility — of who the users are, what applications they're accessing, and when and where they're connecting — is one of the biggest challenges for data center teams. Achieving this gives teams a comprehensive view of devices and connections so they can track applications and users. Application and user insights simplify network security management functions,

such as installation, deployment, and maintenance, from a single console across all locations. This level of visibility also allows for the automation of effective Zero Trust policy implementation and deployment.



TIP

With greater visibility into the data center, network security teams can understand who is accessing what, when, and where, both inside the data center and across multi-cloud environments. Read Chapter 3 to learn more about gaining complete visibility in hybrid cloud environments.

Minimize the attack surface

The best way to protect against advanced attacks, such as advanced persistent threats (APTs) or ransomware, is to minimize opportunities for attack and prevent the lateral movement of any threat in the data center. Segmenting the environment into trust zones helps protect critical applications and shared services from lateral movement attacks by limiting network communication to only the necessary connections. Examples of common trust boundaries include the boundary between the Internet and the data center, a public cloud environment and an on-premises environment, an environment that hosts information governed by regulatory compliance and the rest of the data center, or two applications in a data center.

Physical, virtualized, and containerized next-generation firewalls enable you to define the boundaries between trust zones and easily integrate into the network fabric or third-party switches. Virtualized next-generation firewalls can also be deployed on hypervisors or integrated into software-defined networks (SDNs) and cloud networks to help define trust zones in virtualized and public cloud environments.



TIP

Read Chapter 3 to learn more about minimizing the attack surface in data centers and hybrid cloud environments.

Automate threat protection

While segmentation reduces the attack surface, a multilayered defense is incomplete without the ability to discover threats and malicious activity, block threats in real time, and automatically isolate infected hosts to minimize business disruption as well as prevent data loss. Inserting automated threat protection at trust zone boundaries is a common strategy for protecting hybrid cloud

environments and requires broad coverage with tight integration, including the following functions:

- » **Threat prevention** to block known threats and vulnerability attacks (network and application)
- » **Zero-touch next-generation firewall insertion methods** using scripting and automation frameworks to provide dynamic aspects of security policy management via plug-in management architectures across clouds
- » **Advanced malware analysis** to automatically identify and protect against zero-day exploits
- » **Security analytics** to analyze endpoint and cloud telemetry for automatic detection and response to stealthy threats and insider attacks based on malicious network or host activity
- » **Endpoint security** to protect servers (physical or virtual, on-premises or cloud) from malware, exploits, and ransomware



TIP

Read Chapter 4 to learn more about automating threat protection in the hybrid cloud environments and data centers.

PLANNING A SECURITY STRATEGY FOR YOUR DATA CENTER AND HYBRID CLOUD

Data centers are evolving to allow swift, flexible application delivery, enabling businesses to move faster. Security must be enforced at multiple places to follow workloads everywhere — on the perimeter, network fabric, and host. Implementing best practices will help you better protect data and application workloads as they move dynamically in your physical data centers and across public and private clouds. Use the following steps to help you roll out your hybrid cloud security strategy:

1. Set goals by defining the ideal future state of your data center network, which should include:

- A standardized, scalable design you can replicate and apply consistently across multiple data centers.

- A strategy that uses whitelisting for positive security enforcement and helps you define trust zones to move toward a Zero Trust architecture. Start with tiered applications and proceed from that model with regard to Zero Trust micro-segmentation.
 - Architecture that consistently protects traffic flows initiated by users or from bare metal servers, virtual machines, and containers hosted in on-premises data centers, private clouds, public clouds, or even branches and campuses.
- 2. Work with stakeholders in IT/Support and Security as well as groups that require data center access, such as Engineering and Legal, to develop an access strategy. You'll want to:**
 - Identify users who need access — and the assets they need to access — to define efficient security policy rules by user group.
 - Enforce granular access by implementing segmentation to create trust zones and minimize the opportunities of attack across north-south and east-west traffic.
 - 3. Assess your data center to understand its current state so you can create a plan to reach your desired future state. You should:**
 - Inventory the physical as well as virtual environment and assets, and determine which assets you should protect first.
 - Work with application, network and enterprise architects, as well as business stakeholders, to learn about typical baseline traffic loads and patterns so you understand normal work behavior.
 - 4. Create a data center segmentation strategy to help reduce risk and business impact by preventing hackers from stealing data as well as stopping malware that gains a foothold in your data center from infecting other systems.**

Use firewalls based on flexible form factors to design a granular segmentation strategy for physical and virtual networks to provide visibility into your data center traffic everywhere — at the perimeter, on the network, and on the host.

(continued)

(continued)

5. Plan to use a best practice methodology to inspect all data center traffic, gain complete visibility, reduce the attack surface, and prevent known and unknown threats everywhere in the data center.

6. Phase in best practices over time.

Start by focusing on the most likely threats to your business and network, and protect your most valuable assets first.

Thoughtfully planning a phased, gradual implementation will help you make a smooth, practical transition away from a “hope for the best” security policy to a best practice policy that safely enables applications, users, and content.

- » Understanding your security responsibilities in the cloud
- » Addressing threats to the hybrid cloud

Chapter 2

Security Challenges in Hybrid Clouds

In this chapter, you learn how cloud security and the shared responsibility model, as well as the constantly evolving threat landscape, create new security challenges in hybrid cloud environments.

Cloud Security and The Shared Responsibility Model

As the cloud has become an integral part of enterprise environments, many organizations have been forced into significant compromises with regard to their public and private cloud environments — trading function, visibility, and security, for simplicity, efficiency, and agility. If an application hosted in the cloud isn't available or responsive, network security controls, which all too often introduce delays and outages, are typically

“streamlined” out of the cloud design. Cloud security trade-offs often include:

- » Simplicity or function
- » Efficiency or visibility
- » Agility or security

Cloud computing technologies enable you to evolve your data center from a hardware-centric architecture where applications run on dedicated servers, to a dynamic and automated environment where pools of computing resources are available on-demand, to support application workloads that can be accessed anywhere, anytime, and from any device.

However, many of the features that make cloud computing attractive to organizations are counter to network security best practices. For example:

- » **Cloud computing doesn't mitigate existing network security risks.** The security risks that threaten your network today don't go away when you move to the cloud. In some ways, the security risks you face when moving to the cloud become more significant. Many data center applications use a wide range of ports, rendering traditional security ineffective. Cybercriminals are creating sophisticated port-agnostic attacks that use multiple vectors to compromise their target and then hide in plain sight, using common applications to achieve their objectives.
- » **Separation and segmentation are fundamental to security; the cloud relies on shared resources.** Security best practices dictate that mission-critical applications and data be separated into secure segments, or trust zones, on the network, based on Zero Trust principles (“never trust, always verify”). On a physical network, Zero Trust is relatively straightforward, using firewalls and policies based on application and user identity. In a cloud environment, direct communication between virtual machines (VMs) and containers within a server host occurs constantly — in some cases, across varied levels of trust, making segmentation a real challenge. Mixed levels of trust, combined with a lack of intra-host traffic visibility by virtualized port-based security offerings, may weaken your security posture.

» **Security deployments are process-oriented; cloud computing environments are dynamic.** The creation or modification of your virtual and containerized workloads can often be done in minutes, though the security configuration for this workload may take hours, days, or weeks. Security delays aren't designed to be burdensome; they're the result of a process that is designed to maintain a strong security posture. Policy changes need to be approved, the appropriate firewalls must be identified, and the relevant policy updates determined. In contrast, virtualization and DevOps teams operate in a highly dynamic environment, with workloads being added, removed, and changed rapidly and constantly. The result is a disconnect between security policy and virtualized workload deployment leading to a weakened security posture.

CLOUD-BASED SAAS APPLICATIONS: I CAN'T SEE CLEARLY NOW

Organizations are adopting SaaS-based application services at a breakneck pace. These applications continue to redefine the network perimeter, providing critical functionality and increasing productivity, but at the same time introduce potential new security and data risks if not properly controlled.

In most organizations that use SaaS applications, users are provided access to a specific list of services that the organization has deemed acceptable or suitable for business purposes. However, given the large number of unique SaaS applications that are readily available on the Internet, many users likely aren't strictly complying with such usage policies and are instead using unsanctioned SaaS applications at work. This practice further increases the risk of data leakage to organizations due to the lack of visibility from regular logs or notifications from unauthorized SaaS applications, as well as additional risk of intermeshing users' personal and work emails. In these situations, a user's personal email account may be attacked, and the attacker may then be able to steal data or compromise the user's work email account.

Cloud-based applications and the data that goes with them are increasingly becoming distributed across hybrid cloud environments to improve the agility of the organization and reduce costs. These hybrid cloud environments consist of private clouds (including on-premises and virtualized data centers), public clouds (including infrastructure-as-a-service, or IaaS, and platform-as-a-service, or PaaS), and software-as-a-service (SaaS) applications (see the “Cloud-based SaaS applications: I can’t see clearly now” sidebar in this chapter), each bringing its own unique agility benefits and security issues.

The concern over data exposure has made cloud security a priority. The challenge has become balancing the organization’s need for agility while improving the security of applications and securing the data as it moves among the various clouds. Gaining visibility and preventing attacks that are attempting to exfiltrate data, both from an external location and through a lateral attack, becomes imperative across all locations where the applications and data reside.

A number of groups within an organization can be responsible for cloud security: the network team, cloud architecture team, security team, apps team, compliance team, or the infrastructure team. However, cloud security is also a shared responsibility between the cloud vendor and the organization with each party being responsible for different aspects, depending on the cloud services provided (see Figure 2-1):

» **Private:** Enterprises are responsible for all aspects of security for the cloud as it is hosted within their own data centers. This includes the physical network, infrastructure, hypervisor, virtual network, operating systems, firewalls, service configuration, identity and access management, and so on. The enterprise also owns the data and the security of the data.

» **Public:** In public clouds, like Google Cloud Platform (GCP), Amazon Web Services (AWS), or Microsoft Azure, the cloud vendor is responsible for the security of the infrastructure, physical network, and hypervisor. The enterprise is responsible for the security of the workload operating system, apps, virtual network, access to the tenant environment or account, and the data.

» **SaaS:** SaaS vendors are primarily responsible for the security of their platform, which includes physical security, infrastructure, and application security. These vendors do not own the enterprise data or assume responsibility for how enterprises use the applications. As such, the enterprise is responsible for security that would prevent and minimize the risk of malicious data exfiltration, accidental exposure, or malware insertion.




FIGURE 2-1: Shared security responsibility across private, public, and SaaS cloud services.



REMEMBER

Cloud security refers to the efforts of securing data, applications, and infrastructure intrinsic to the use of cloud computing — including policies, technologies, and controls.

As a company expands its digital footprint across the hybrid cloud, including private and public cloud environments and SaaS applications, the responsibility for securing data, applications, and infrastructure changes, but regardless of the platform used, the enterprise will always be responsible for ensuring the security of its own data.

In order to safely enable applications and their associated data, IT security must be confident that its cloud vendors have implemented the appropriate security measures to keep your applications and data secure. To compensate for what cloud vendors do not secure, an organization must also have the right tools in place to manage and secure the risks effectively in order to keep their data secure. These tools must provide visibility into activity within SaaS applications, detailed analytics on cloud and SaaS usage to prevent data risk and compliance violations, context-aware policy controls to drive enforcement and quarantine cloud workloads if

a violation occurs, and real-time threat intelligence and detection on known and unknown threats to prevent new malware insertion points.



REMEMBER

Cloud security is a shared responsibility between the cloud vendor and the organization. However, the organization is always responsible for securing its own data.

Here are some key requirements for securing a hybrid cloud environment:

- » **Consistent security in physical and virtualized form factors.** The same levels of application control and threat prevention should be used to protect both your cloud computing environment and your physical network. First, you need to be able to confirm the identity of your applications, validating their identity and forcing them to use only their standard ports. You also need to be able to block the use of rogue applications while simultaneously looking for and blocking misconfigured applications. Finally, application-specific threat prevention policies should be applied to block both known and unknown malware from moving into and across your network and cloud environment.
- » **Segment your business applications using Zero Trust principles.** In order to fully maximize the use of computing resources, it's now a relatively common practice to mix application workload trust levels on the same compute resource. Although efficient in practice, mixed levels of trust introduce new security risks in the event of a compromise. Your cloud security solution needs to be able to implement security policies to create trust zones based on the concept of Zero Trust, as a means of controlling traffic between workloads while preventing lateral movement of threats.
- » **Centrally manage security deployments; streamline policy updates.** Physical network security hardware is still deployed in almost every data center, so it's critical that you have the ability to manage both hardware and virtual form factor deployments from a centralized location using the same management infrastructure and interface. In order to ensure that security keeps pace with the speed of change your workflows may exhibit, your security solution should include automation features that will allow you to lessen, and in some cases, eliminate the manual processes that security policy updates often require.



WARNING

Many cloud security offerings are merely virtualized versions of port- and protocol-based security appliances, delivering the same inadequacies as their physical counterparts. Major business requirements for cloud security include:

- » Preventing threats
- » Scalability and automation
- » Keeping pace with the business

Preventing threats has become more difficult in the past several years. Basic attacks on infrastructure have given way to multi-vector, application-borne, sophisticated attacks that are stealthy, profit-driven, unwittingly aided by enterprise users, and in many cases, polymorphic. The level of organization associated with the development of these threats is also unprecedented.

Regulatory and compliance requirements — such as the Payment Card Industry's (PCI) Data Security Standards (DSS), U.S. healthcare mandates like the Health Insurance Portability and Accountability Act (HIPAA), and privacy regulations like the European Union (EU) General Data Protection Regulation (GDPR), Australian Privacy Principles, and the California Consumer Privacy Act (CCPA) — are pushing network segmentation deeper into organizations generally, and into data centers and cloud environments specifically.

Finally, needless complexity can introduce integration issues, outages, and latency. Keeping the data center and hybrid cloud design and architecture as consistent as possible is essential to improving performance, availability, manageability, and security.

The Dynamic Nature of Modern Threats

The modern threat landscape is constantly evolving, and many sophisticated new threats have emerged in recent years. Email and web browsers are still the main attack vectors today, with malicious content either attached or downloaded as an executable or macro-based file. The malicious use of remote access applications is another significant attack vector. Threats that directly target applications can pass right through the majority of enterprise defenses, which have historically been built to

provide network-layer protection. Threat developers exploit various methods to infiltrate networks, including:

- » **Port hopping**, where ports/protocols are randomly shifted over the course of a session
- » **Use of nonstandard ports**, such as running Yahoo! Messenger over Transmission Control Protocol (TCP) port 80 (HyperText Transfer Protocol, or HTTP) instead of the standard TCP port for Yahoo! Messenger (5050)
- » **Tunneling within commonly used services**, such as when sharing files or using messaging applications like Telegram Messenger
- » **Hiding within Secure Sockets Layer (SSL) encryption**, which masks the application traffic, for example, over TCP port 443 (HyperText Transfer Protocol Secure, or HTTPS)



WARNING

Legitimate applications are being used by attackers to spread malware.

The evasion techniques built into these and many other modern applications are being leveraged to provide threats with “free passage” into enterprise networks. So, it’s no surprise that more than 80 percent of all new malware and intrusion attempts are exploiting weaknesses in applications, as opposed to weaknesses in networking components and services. Together with the implicit trust that users place in their applications, all these factors combine to create a “perfect storm.” The motivation for attackers has also shifted — from gaining notoriety to political activism, espionage, and making money. The name of the game today is information theft. Consequently, it’s no longer in an attacker’s best interests to devise threats that are “noisy” or that are relatively benign. To be successful, a thief must be fast or stealthy — or both.

For those attackers who favor speed over sophistication — speed of initial threat generation, speed of modification, and speed of propagation — the goal is to develop, launch, and quickly spread new threats immediately on the heels of the disclosure of a new vulnerability. The resulting zero-day and near-zero-day exploits then have an increased likelihood of success because reactive countermeasures, such as patching and those tools that rely on threat signatures (such as antivirus software and intrusion prevention), have trouble keeping up — at least during the early phases of a new attack.

This speed-based approach is facilitated in large part by the wide-spread availability of threat development websites, toolkits, and frameworks. Unfortunately, another by-product of these resources is the ability to easily and rapidly convert “known” threats into “unknown” threats — at least from the perspective of signature-based countermeasures. This transformation can be accomplished either by making a minor tweak to the code of a threat, or by adding entirely new propagation and exploit mechanisms, thereby creating what is commonly referred to as a *blended threat*.

Many of today’s threats are built to run covertly on networks and systems, quietly collecting sensitive or personal data, and going undetected for as long as possible. This approach helps to preserve the value of the stolen data and enables repeated use of the same exploits and attack vectors. As a result, threats have become increasingly sophisticated. Rootkits, for example, have become more prevalent. These kernel-level exploits effectively mask the presence of other types of malware, enabling them to persistently pursue the nefarious tasks they were designed to accomplish (such as intercepting keystrokes).

Encryption is increasingly used to secure not just sensitive or private information, but practically all traffic traversing enterprise networks. However, organizations are essentially left blind to any security threats contained inside encrypted traffic. Attackers exploit this lack of visibility and identification to hide within encrypted traffic and spread malware. Even legitimate websites that use SSL can be infected with malware. Moreover, attackers increasingly use SaaS applications to deliver malware. For example, an attacker can place a malicious file on a website with encryption and host a file to be downloaded.



WARNING

Without the ability to decrypt, classify, control, and scan SSL-encrypted traffic, it’s impossible for an organization to adequately protect its business and its valuable data from modern threats.

Threats to enterprise data center and hybrid cloud environments include:

- » Ransomware
- » Credential theft
- » Domain Name System (DNS) based attacks
- » Targeted “low-and-slow” attacks and advanced persistent threats (APTs)

Ransomware

Ransomware has existed in various forms for decades, but in the last few years, criminals have perfected the key components for these types of attacks. Ransomware uses malware to encrypt a victim's data until a ransom is paid — usually in cryptocurrency. Ransomware has become a multimillion-dollar criminal business targeting both individuals and corporations (see the “Ransomware: LockerGoga” sidebar in this chapter). Because of its low barriers to entry and effectiveness in extorting ransom payments from its victims, the spread of ransomware has increased exponentially in recent years. A typical ransomware attack consists of the following steps:

- 1. Compromise and control a system or device.**

Most ransomware attacks begin by using social engineering to trick users into opening an attachment or viewing a malicious link in their web browser. This allows attackers to install malware onto a system and take control.

- 2. Prevent access to the system.**

Attackers will either identify and encrypt certain file types or deny access to the entire system.

- 3. Notify the victim.**

Though seemingly obvious, attackers and victims often speak different languages and have varying levels of technical capabilities. Attackers must alert the victim about the compromise, state the demanded ransom amount, and explain the steps for regaining access.

- 4. Accept ransom payment.**

To receive payment while evading law enforcement, attackers utilize cryptocurrencies such as Bitcoin for the transaction.

- 5. Restore full access (usually).**

Attackers must return access to the device(s). Failure to restore the compromised system(s) destroys the effectiveness of the scheme — no one would be willing to pay a ransom if they didn't believe access to their data would be restored.



WARNING

Although an attacker usually restores access to the victim's data after the ransom is paid, there is no “money-back guarantee.” There's also no guarantee that the attacker didn't also steal a copy of your data and sell it on the dark web.

RANSOMWARE: LOCKERGOGA

The LockerGoga ransomware was first publicly reported in January 2019 by Bleeping Computer, which tied the malware to an attack against French engineering company Altran Technologies. Several variants have since been found in the wild, where they were used in attacks against Norwegian aluminum manufacturer Norsk Hydro and two chemical companies: Hexion and Momentive.

Currently, LockerGoga does not support any worm-like capabilities that would allow it to self-propagate by infecting additional hosts on a target network. LockerGoga has been observed moving around a network via the Server Message Block (SMB) protocol, which indicates the actors simply manually copy files from computer to computer.

LockerGoga's developers continue to add capabilities and launch new attacks. The addition of WS2_32.dll and use of undocumented Windows application programming interface (API) calls indicates a level of sophistication beyond typical ransomware authors. The former could lead to the eventual inclusion of command-and-control (C2) communication or automated propagation, and the latter requires some working knowledge of Windows internals.

These features raise more questions about the actor's intent, as ransomware is typically one of the least advanced forms of malware. Are they motivated by profits or something else? Has the motive changed over time? Why would developers put so much effort into their work only to partially encrypt files? Why do they include an email address, rather than seeking payment through more frequently used cryptocurrencies?



REMEMBER

Command-and-control (C2) refers to 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.

Credential theft

Users and their credentials are among the weakest links in an organization's security infrastructure. According to Forrester Research, at least 80 percent of data breaches today involve

compromised credentials. Credential theft has become so prevalent in the attackers' playbook that it's often said that attackers no longer hack into a target network — they simply log in. The primary techniques that attackers use to steal credentials include:

- » Social engineering
- » Phishing and malware
- » Brute force
- » Security question reuse
- » Reusing stolen passwords or shared credentials sold on the dark web

Attackers use these credentials to gain access to a network, move laterally, and escalate their privileges for unauthorized access to applications and data (see the “Credential theft: Shamoon 2” sidebar in this chapter).

With stolen credentials as part of their toolset, attackers' chances of successfully breaching go up, and their risk of getting caught goes down. To prevent credential theft, most organizations rely on employee education, which is prone to human error by nature. Technology products commonly rely on identifying known phishing sites and filtering email. However, these methods can sometimes be bypassed — checking for known bad sites misses newly created ones, and attackers can evade mail filtering technology by sending links through social media.



TIP

Organizations should look for a firewall with machine learning-based analysis to identify websites that steal credentials. If the analysis identifies a site as malicious, the firewall should be automatically updated in real time and block it. Still, there will always be new, never-before-seen phishing sites that are treated as “unknown.” Your firewall must allow you to block submission of user credentials to unknown sites. The firewall must also allow you to protect sensitive data and applications by enforcing multi-factor authentication (MFA) to prevent attackers from abusing stolen credentials. By integrating with common MFA vendors, your firewall can protect your applications containing sensitive data, including legacy applications.

CREDENTIAL THEFT: SHAMOON 2

Palo Alto Networks Unit 42 researchers have been following the Shamoon 2 attacks closely since November 2016. Credential theft is a key part of Shamoon 2 attacks. Shamoon 2 enters and spreads through an organization in three stages:

1. Shamoon 2 attackers access and compromise a single system in the network using Remote Desktop Protocol (RDP) with stolen, legitimate credentials. This becomes their distribution server; they download their tools and malware to this system.
2. Attackers execute commands on the distribution server to connect to specific, named systems on the network using the stolen, legitimate credentials, and infect them with the Disttrack malware.
3. The Disttrack malware executes on those named systems the attacker has successfully infected. The Disttrack malware attempts to connect to and spread itself to up to 256 IP addresses on its local network. Any systems successfully infected in this stage also attempt to infect up to 256 IP addresses on their local networks.

Shamoon 2 attacks are targeted to a specific region, but it would be a mistake to disregard the threat. Shamoon 2 attackers are using a rudimentary, but effective, distribution system of their own making. The power of their attack doesn't lie in the tools they use, but in their ability to obtain and abuse legitimate credentials.

DNS-based attacks

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



At a high level, a DNS lookup generally involves the following steps (see Figure 2-2):

1. A user initiates a query by typing a Uniform Resource Locator (URL) into a web browser. The query is sent to a DNS resolver, which is a server usually provided by an IP provider. The DNS resolver matches the URL to its corresponding IP address.
2. The resolver queries one of the DNS root servers. Root servers are distributed around the world and hold the locations of all top-level domains (TLDs) such as .com, .edu, and .net.
3. The root server sends a response to the resolver.
4. The resolver can now query the appropriate TLD server.
5. The TLD server doesn't hold the IP address for specific domains, but it knows the locations of the authoritative name servers for specific domains.
6. The DNS resolver queries one of the authoritative name servers.
7. The authoritative name server knows the IP address and responds with an address record.
8. The DNS resolver sends the IP address of the web server back to the user's device.
9. The website loads in the browser.




FIGURE 2-2: The DNS resolution process.

All this happens in the background within a few milliseconds. Sites like google.com or paloaltonetworks.com may have multiple IP addresses, which can speed up DNS lookup times. Millions of people, even from different countries around the world, may be looking for the same information at the same time. These queries will likely go to different servers that are distributed worldwide.

DNS information is also cached on your computer and on the servers used by your Internet service provider. Once the IP address for a particular URL is saved, your computer no longer needs to access a DNS resolver to resolve the name with its IP address.

DNS is an open service, and by default it does not have a way to detect DNS-based threats. As a result, malicious activity within DNS can be used to propagate an attack causing costly damage and downtime.

DNS is a massive and often overlooked attack surface, present in every organization, that can be used for malware delivery, C2 communications, and data exfiltration. Adversaries take advantage of the open and widespread nature of DNS to abuse it at different communication points during the back-and-forth DNS resolution process described in the preceding steps.

According to the Palo Alto Networks Unit 42 threat research team, almost 80 percent of malware uses DNS to initiate C2 communications (see the “DNS-based attacks: OilRig” sidebar in this chapter). Attackers establish reliable command channels that are difficult to take down or identify since DNS is such a reliable way to maintain a connection to DNS servers. Once a connection is established, attackers can use DNS traffic to deliver malware into a network or tunnel data out.

Unfortunately, security teams often lack basic visibility into how threats use DNS to establish and maintain control of infected devices. Adversaries take advantage of the ubiquitous nature of DNS to abuse it at multiple points of an attack, including reliable C2. Security teams also struggle to keep up with new malicious domains and enforce consistent protections for millions of emerging domains at once. Attackers develop domain generation algorithms (DGAs), which automatically create thousands of malicious domains that can be used for C2. As adversaries

increasingly automate their attacks, it becomes almost impossible to identify and stop these threats. It's impossible for enterprise network and security teams to keep up with the high volume of malicious domains, let alone advanced tactics like DNS tunneling for stealthy data theft and DNS hijacking to redirect legitimate DNS queries to malicious sites.



TIP

You cannot simply blacklist attacks that use DNS as this tactic often relies on relatively static threat feeds that work off known bad domains. Without analytics, it is impossible to predict highly dynamic malicious domains. Stopping attacks that use DNS requires a next-generation firewall that can apply predictive analytics and machine learning to identify unknown bad domains dynamically.

DNS-BASED ATTACKS: OILRIG

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

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

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

Targeted “Low-and-Slow” Attacks and APTs

Targeted attacks and APTs against specific organizations or individuals are another major concern. In this case, attackers often develop customized attack mechanisms to take advantage of the specific equipment, systems, applications, configurations, and even personnel employed in a specific organization or at a given location, and quietly collect sensitive data over extended periods. These “low-and-slow” tactics are designed to avoid detection for as long as possible. Whereas the average time for an organization to identify a breach in 2019 was 206 days according to the Ponemon Institute, a targeted “low-and-slow” attack or APT may go undetected for years (see the “Carbanak: The great bank robbery” sidebar in this chapter).

APTs are a class of threats that often combine advanced malware and botnet components to execute a far more deliberate and potentially devastating attack than other types of attacks. As the name applies, an APT has three defining characteristics:

- » **Advanced:** In addition to advanced malware and botnets, the attackers typically have the skills to develop additional exploitation tools and techniques and may have access to sophisticated electronic surveillance equipment, satellite imagery, and even human intelligence assets.
- » **Persistent:** An APT may persist over a period of many years. The attackers pursue specific objectives and use a low-and-slow approach to avoid detection. The attackers are well organized and typically have access to substantial financial backing to fund their activities, such as a nation-state or organized crime.
- » **Threat:** An APT is a deliberate and focused, rather than opportunistic, threat that can cause real damage.



REMEMBER

A *botnet* is a broad network of malware-infected endpoints (bots) working together and controlled by an attacker through C2 infrastructure.

The increasing speed and sophistication of threats emphasize the need for proactive countermeasures with extensive visibility and control at the application layer of the network computing stack.

CARBANAK: THE GREAT BANK ROBBERY

Carbanak is one example of a targeted attack that began in August 2013 and is currently still active despite the arrest of the alleged Carbanak “mastermind” in March 2018. The attackers have sent spear-phishing emails with malicious Control Panel file (.cpl) attachments or Word documents exploiting known vulnerabilities. When an initial system has been compromised, additional reconnaissance is performed to identify automated teller machines (ATMs), financial accounts, or other areas where money can be transferred for eventual extraction. Each raid has lasted two to four months. To date, the attackers have targeted more than 100 financial institutions and businesses, causing aggregated losses estimated at more than \$1 billion.

IN THIS CHAPTER

- » Seeing your entire protect surface
- » Segmenting data centers and hybrid cloud environments
- » Migrating security policies dynamically with hybrid cloud workloads

Chapter 3

Delivering Consistent Security Using Zero Trust

Conventional security models operate on the outdated assumption that you can trust everything inside your network. However, given the increased sophistication of attacks and insider threats, you need new security measures to stop them from spreading once they're inside. Because traditional security models are designed to protect your perimeter, threats that get inside your network often go undetected and are free to morph and move wherever they choose to extract sensitive business data. In the digital world, trust is a vulnerability.

A Zero Trust approach to security helps prevent data breaches by eliminating the concept of inherent trust in an organization's data center or hybrid cloud environment. Rooted in the principle of "never trust, always verify." The key to adopting a Zero Trust approach in data centers and hybrid cloud environments lies in leveraging a multi-layered network security posture that delivers consistent network visibility, segmentation and micro-segmentation, and threat detection and response across physical, virtual, and containerized infrastructures. Preventing lateral movement relies on security that is context-based, dynamic, and integrated throughout your organization.

In this chapter, you learn how Zero Trust helps you ensure an effective and consistent security posture across your data center and hybrid cloud environment.

Gaining Complete Visibility

A Zero Trust architecture requires complete visibility of your organization's protect surface, which is orders of magnitude smaller than its attack surface and is always knowable. Your organization's protect surface consists of all your critical data, applications, assets, and services (DAAS). This can include:

- » **Data** such as protected health information (PHI), personally identifiable information (PII), financial information, and intellectual property (IP)
- » **Applications** including commercial off-the-shelf software, software-as-a-service (SaaS), and custom developed software
- » **Assets** such as supervisory control and data acquisition (SCADA) systems, point-of-sale (POS) terminals, medical equipment, manufacturing systems, and Internet of Things (IoT) devices
- » **Services** such as Domain Name System (DNS), Dynamic Host Configuration Protocol (DHCP), Active Directory, and Lightweight Directory Access Protocol (LDAP)

As discussed in Chapter 1, legacy port-based firewalls do a poor job of identifying applications, content, and users. A next-generation firewall enables comprehensive Layer 7 visibility of your entire data center and hybrid cloud protect surface (see Figure 3-1). It performs true classification of data and application traffic, based not simply on port and protocol (like a port-based firewall) but on contextual factors such as a user, their device, and the applications they need to use to perform their role throughout the day. This classification and filtering activity occurs as an ongoing process of application analysis, decryption, decoding, and heuristics as well. These capabilities progressively peel back the layers of a traffic stream to determine its true application identity.

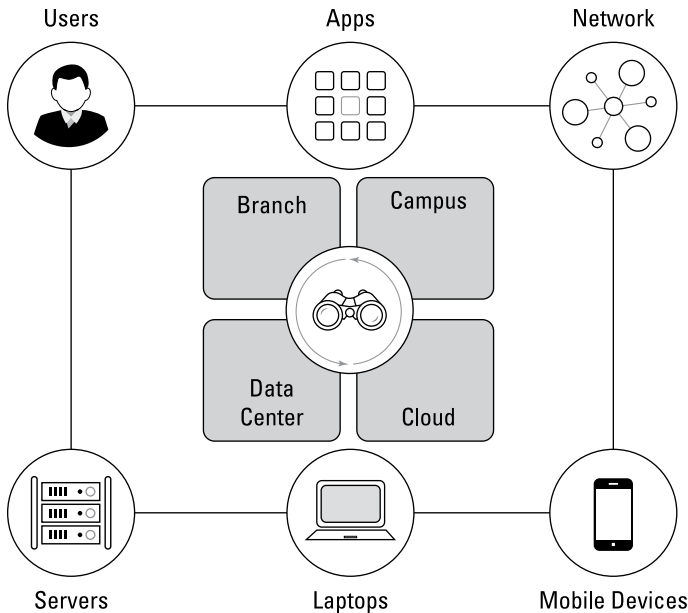


FIGURE 3-1: Deep visibility into user, device, network, and application activity is essential to hybrid cloud and data center security.

Three core capabilities in a next-generation firewall that enable complete visibility for a Zero Trust architecture include:

- » Application identification
- » User identification
- » Content identification

Application identification

The first step in application identification is to establish the port and protocol. Next, robust application identification and inspection enables granular control of the flow of sessions through a firewall based on the applications that are being used.

Most enterprise network traffic is now encrypted, and attackers exploit encryption to hide threats from security devices. This means even businesses with mature, comprehensive security measures in place can be breached if they aren't monitoring encrypted traffic.

The ability to decrypt Secure Sockets Layer (SSL) and other encrypted traffic, such as Secure Shell (SSH), is a foundational security function of next-generation firewalls. Key capabilities include recognition and decryption on any port (inbound or outbound), policy control over decryption, and the necessary elements to perform decryption across tens of thousands of simultaneous SSL connections with predictable performance.

Positive application identification is the traffic classification engine at the heart of next-generation firewalls. It requires a multifaceted approach to determine the identity of applications on the network, regardless of port, protocol, encryption, or evasive tactics. Application identification techniques used in next-generation firewalls include:

»» **Application protocol detection and decryption:**

Determines the application protocol and if encrypted, it decrypts the traffic so that it can be analyzed further. Traffic is re-encrypted after all the next-generation technologies have an opportunity to operate.

»» **Application signatures:** Context-based signatures look for unique properties and transaction characteristics to correctly identify the application regardless of the port and protocol being used. This includes the ability to detect specific functions within applications, such as file transfers within software-as-a-service (SaaS) applications.

»» **Heuristics:** For traffic that eludes identification by signature analysis, heuristic (or behavioral) analyses are applied — enabling identification of any suspicious applications, such as peer-to-peer (P2P) or Voice over Internet Protocol (VoIP) tools that use proprietary encryption.



TIP

Having the technology to accurately identify applications is important but understanding the security implications of an application so that informed policy decisions can be made is equally important. Look for next-generation firewalls that include information about each application, and its behaviors and risks, to provide IT administrators with application knowledge such as known vulnerabilities, ability to evade detection, file transfer capabilities, bandwidth consumption, malware transmission, and potential for misuse.

User identification

User identification technology links Internet Protocol (IP) addresses to specific user identities, enabling visibility and control of network activity on a per-user basis. Tight integration with Lightweight Directory Access Protocol (LDAP) directories, such as Microsoft Active Directory (AD), supports this objective in two ways:

- » It regularly verifies and maintains the user-to-IP address relationship using a combination of login monitoring, end-station polling, and captive portal techniques.
- » It communicates with AD to harvest relevant user information, such as role and group assignments.

These details are then available to:

- » Gain visibility into who specifically is responsible for all application, content, and threat traffic on the network, including users on mobile devices, working remotely, or located in branch offices
- » Enable the use of identity as a variable within access control policies
- » Facilitate troubleshooting/incident response and reporting

User identification is also an important capability to help prevent credential theft and abuse. The majority of network breaches today involve stolen credentials that attackers use to simply log on to the network (rather than hacking in) and elevate privileges leveraging other stolen credentials once inside the network.

With user identification, IT departments get another powerful mechanism to help control the use of applications in an intelligent manner. For example, a remote access application that would otherwise be blocked because of its risky nature can be enabled for individuals or groups that have a legitimate need to use it, such as IT administrators.

Content identification

Content identification infuses next-generation firewalls with capabilities previously unheard of in enterprise firewalls, including:

- » **Threat prevention:** This component prevents malware and exploits from penetrating the network, regardless of the application traffic in which they are hiding.
 - *Application decoder:* Pre-processes data streams and inspects for specific threat identifiers.
 - *Stream-based malware scanning:* Scanning traffic as soon as the first packets of a file are received — as opposed to waiting until the entire file is in memory — maximizes throughput and minimizes latency.
 - *Uniform threat signature format:* Performance is enhanced by avoiding the need to use separate scanning engines for each type of threat. Viruses, command-and-control (C2) communications, and vulnerability exploits can all be detected in a single pass.
 - *Vulnerability attack protection:* Similar to the functionality provided in intrusion prevention systems (IPS), protocol anomaly, behavior anomaly, and heuristic detection mechanisms are used for protection from known and unknown threats.
 - *Cloud-based intelligence:* For content that's unknown, the ability to send to a cloud-based security service ("sandboxing") for rapid analysis and a "verdict" that the firewall can then use.
- » **Uniform Resource Locator (URL) filtering:** URL filtering is a tool used to classify content. An integrated URL database allows administrators to monitor and control web surfing activities of employees and guest users. Employed in conjunction with user identification, web usage policies can even be set on a per-user basis, further safeguarding the enterprise from an array of legal, regulatory, and productivity-related risks.
- » **File and data filtering:** Taking advantage of in-depth application inspection, file and data filtering enables enforcement of policies that reduce the risk of unauthorized information transfer, or malware propagation. Capabilities

include the ability to block files by their actual type (not based on just their extension), and the ability to control the transfer of sensitive data patterns like credit card numbers. Granular policies enable organizations to bypass decryption of certain sensitive data, such as data to and from a known financial institution, if required by security and/or privacy compliance mandates. This complements the granularity of application identification, which offers the ability to control file transfer within an individual application.

With content identification, IT departments gain the ability to stop threats, reduce inappropriate use of the Internet, and help prevent data leaks — all without having to invest in a pile of additional threat prevention products that cause appliance sprawl, don't work well because of their lack of integration, and lack comprehensive visibility.

Minimizing Your Attack Surface with Segmentation

A flat unsegmented network is difficult to protect because if an attacker gains access to the network, the attacker can easily move laterally and compromise critical systems. This is particularly true in a hybrid cloud environment where on-premises data centers are connected to private and public clouds. Old segmentation methods such as virtual local area networks (VLANs) don't scale well, are difficult to automate, and don't take into account users, content, or applications, so they provide little control over or visibility into traffic.

Creating a segmentation strategy that provides granular access control to hybrid cloud resources will give you better control over traffic. The more granular your segmentation strategy, the more control over traffic you gain because traffic must traverse a firewall as it flows between segments, or trust zones. Many organizations use the application itself as the trust boundary, essentially putting each application in its own trust zone, protected by a next-generation firewall to inspect all traffic that traverses the boundary. For the most critical applications, you may also want to leverage a micro-segmentation tool to restrict traffic moving between workloads within a trust zone, to further reduce

risk. Segmentation also makes compliance and compliance audits easier because you can prevent all but the necessary access to personal information, which protects the data and reduces the scope of audits.

In a hybrid cloud environment, there are two different types of traffic, each of which is secured in a different manner (see Figure 3-2):

- » **North-south** refers to data packets that move in and out of, as well as between data center, cloud, and WAN environments. North-south traffic is secured by one or more perimeter edge firewalls that control all the traffic into and out of the data center. In on-premises data centers, this is typically a physical firewall, whereas in public cloud environments, a virtual firewall is used.
- » **East-west** refers to data packets moving between virtual machines, containers, and application workloads entirely within the data center or cloud environment. East-west traffic is protected by virtualized firewalls instantiated on hypervisors or within container clusters. East-west firewalls are inserted transparently into the application infrastructure, often positioned closest to the actual virtualized workload, and do not necessitate a redesign of the logical topology.

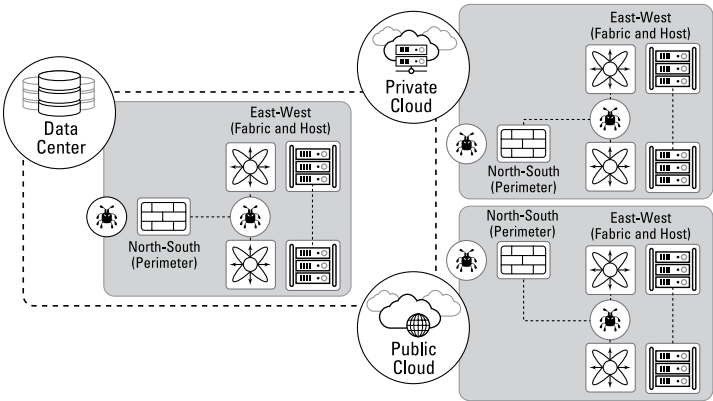


FIGURE 3-2: Multi-layered segmentation restricts lateral movement across hybrid cloud environments by attackers.

Historically, organizations would implement security to protect traffic flowing north–south, which is insufficient for protecting east–west traffic which now constitutes the majority of data center (including private and public cloud) traffic.



TIP

To improve their security postures with regard to sensitive data, organizations recognize that protecting against threats across the entire computing environment, both north–south and east–west has rapidly become a security best practice.




TIP

A next–generation firewall deployed at the trust boundary between trust zones enables a Zero Trust architecture that limits the scope of an attack and blocks lateral movement using a combination of segmentation and threat protection.

Using Dynamic Security to Support Moves and Changes

Today’s data centers and hybrid cloud environments move and change rapidly. Virtualized and containerized workloads are created and deprecated continually as applications are built and modified. It’s imperative that security policy provisioning and updates in these environments keep up with the pace of application development and change.

Traditionally, security policies are applied to workloads based on IP addresses. In hybrid cloud environments, IP addresses are often dynamic in nature, so it becomes essential to abstract these addresses away from the security policies themselves. Tags are the main mechanism by which next–generation firewalls overcome this challenge. A tag creates a grouping of IP addresses that can be used to formulate a policy. For example, a tag can be created for “map–servers” and a security policy can be created that uses the “map–servers” tag as the source address. As map–server workloads are created and deprecated, they can dynamically enter and leave the map–server group, eliminating the need to update the policy with each change, while still ensuring that any workload tagged appropriately will be secured (see Figure 3–3).

		Source		
Name	Tag	Zone	Address	User
SSH rule	none	any	 map-servers	any

An empty bucket into which IP addresses are populated



FIGURE 3-3: Tags abstract policy away from specific IP addresses.

Tag-based policies also enable the flexibility to apply different rules to the same server based on tags that define its role on the network, the operating system, or the different kinds of traffic it processes. Tags can even be used to apply specific policies to workloads based on regulatory compliance standards.



TIP

Because tags are applied to workloads as part of the infrastructure provisioning process, it's important to agree on a unified system for tagging across your organization. Key stakeholders from infrastructure, security, networking, and application teams should all agree to a tagging methodology and nomenclature to avoid confusion and miscommunication. Tagging methodologies can also carry over to public cloud resources.



TIP

THE TRUTH ABOUT ZERO TRUST

As Zero Trust has become more widely known, so too have the misconceptions around what Zero Trust is and how to achieve a Zero Trust network architecture. Here are four prevalent myths about Zero Trust and the truths behind them.

Myth 1: The goal of Zero Trust is to make a system trusted.

Truth: The goal of Zero Trust is to eliminate the concept of inherent trust so that you can strategically protect what's important to your organization.

Myth 2: Zero Trust is complex, costly, and time consuming.

Truth: Start by focusing on the most critical applications and data sets. Build your strategy around the four design concepts of Zero Trust:

- Define business outcomes
- Design from the inside out
- Determine who or what needs access
- Inspect and log all traffic

Myth 3: Zero Trust is all about identity.

Truth: Identity is only part of Zero Trust. Traffic that the asserted identity generates must be inspected for malicious content and unauthorized activity, and logged through Layer 7 (the Application Layer). Start with the users and data in your organization's protect surface, then extend across the network to the applications, assets, and services (DAAS: data, applications, assets, and services).

Myth 4: You can do Zero Trust at Layer 3 (the Network Layer).

Truth: Most attackers can easily bypass traditional network firewalls operating at Layers 3 and 4 (the Transport Layer) using port scans to identify vulnerable open ports or services. When you create policy at Layer 7, you have visibility throughout the entire stack, preventing attackers from moving across the internal network and accessing sensitive data or systems.

IN THIS CHAPTER

- » Revisiting defense in depth
- » Redefining the security perimeter
- » Exploring key components of threat protection
- » Working together to deliver effective threat protection

Chapter 4

Leveraging Unmatched Threat Protection

In this chapter, you explore the challenges of traditional defense in depth and perimeter-based security best practices, learn the essentials of threat protection, and discover why a broad ecosystem of security partners is critical to delivering unmatched threat protection in data centers and hybrid cloud environments.

The Challenges of Defense in Depth

Defense in depth is a well-understood and long-accepted security maxim. The basic idea of defense in depth is to deploy multiple security detection and prevention solutions to address a wide array of potential threats. For example, an organization might typically deploy firewalls, intrusion preventions (IPS), and anti-virus software. Over time, a whole litany of new point security solutions (“innovations”) from different security vendors were developed and sold, including web content filtering, email security (anti-spam) gateways, web application firewalls (WAFs), data loss prevention (DLP), distributed denial-of-service (DDoS) defense, network vulnerability scanning, cloud access security brokers (CASB), security information and event management (SIEM), and much more.

As a result, defense in depth, for many, has become defense ad nauseam. Unfortunately, there are very real consequences of a poorly planned defense in depth strategy. Aside from the obvious, and likely exorbitant, cost of purchasing and maintaining multiple point security solutions for a hybrid cloud environment, there is the inherent complexity that comes with deploying and operating numerous point security solutions. Complexity creates numerous security challenges including:

- » **Lack of a cohesive, end-to-end security strategy.** Many point security solutions have overlapping capabilities and functions. Inevitably, there will also be capability and coverage gaps. Without a complete understanding of the capabilities and functionality in all the different solutions and a coherent cybersecurity policy, the enterprise security posture will be weakened.
- » **Limited integration and interoperability.** Not all point security solutions play well with others. In fact, most don't. This makes it difficult for security teams to correlate events and respond to threats in real time. It can also lead to siloed operations and finger-pointing during troubleshooting and incident response. So now, not only do your toys not play well together — neither do your network and security teams!
- » **Error-prone configuration and operation.** Network and security professionals must learn how to properly configure and operate each of these point security solutions. They all have different operating systems, dashboards, command syntaxes, and more, even if they are all made by the same security vendor — which they rarely are. As a result, mistakes are made, and security vulnerabilities are unwittingly exposed.
- » **Shortage of security skills.** The worldwide shortage of security professionals is an ongoing problem for every enterprise today and for the foreseeable future. Attaining, training, and retaining top security talent is hard enough. Add to that the frustration that comes with having to settle for being a “jack of all trades, master of none” and trying to sleep at night worrying about becoming the next big data breach story and it's easy to understand why security talent is in short supply.



REMEMBER

Complexity is the enemy of an effective enterprise cybersecurity strategy.

What Is a Perimeter?

In the not-too-distant past, security vendors, network architects, and security practitioners described networks in terms of the “untrusted” public Internet and the “trusted” internal corporate network with firewalls deployed at the network perimeter. But the network perimeter has become a relic of a bygone era when everything was simple: black or white, good or bad, “trusted” or “untrusted.”

The reality is that attackers have always exploited relatively weak security designs that relied on the firewall as the arbiter of trust between the Internet and the corporate network. Once inside, attackers had — and continue to have — free rein in the data center and on the network, because trust is assumed. This threat is further exacerbated by the fact that legacy port-based firewalls deployed at the network perimeter only inspect north-south traffic (traffic passing between different zones, such as from an on-premises data center to the Internet or to a public cloud). These firewalls have no visibility into east-west traffic (traffic between systems and applications inside the data center or cloud), which today constitutes the majority of data center and hybrid cloud network traffic.

Further exacerbating the challenges of a perimeter-based security architecture is the fact that most enterprises today operate a hybrid cloud environment composed of a combination of on-premises data centers, private clouds, and public clouds including software-as-a-service (SaaS), platform-as-a-service (PaaS), and infrastructure-as-a-service (IaaS) offerings. Mobile and remote users are also accessing enterprise computing resources in data centers and in the cloud from a multitude of devices from practically anywhere in the world. Thus, the traditional network perimeter today is everywhere — and yet it is nowhere.

Today, perimeter-based security has to be defined at a more granular level than, say the logical boundary between a data center environment and the Internet. Perimeters, or trust boundaries between trust zones, must be defined at multiple layers within a hybrid cloud for both north-south and east-west traffic. A trust zone may be composed of a single resource, such as a virtual machine (VM), or a group of resources within a defined virtual network. Trust zones and trust boundaries are also dynamic. They

may move automatically from an on-premises data center to a public cloud to provide business continuity in the event of an outage, or they may expand on-demand from a single VM (or container) to dozens of VMs or hundreds of containers in a VM scale set (or in a Kubernetes cluster) during periods of peak demand.

Thus, modern security solutions for hybrid clouds must be able to move and scale dynamically with the resources they are deployed to protect. They cannot be constrained to physical or logical boundaries, such as an on-premises data center or IP subnet, respectively.

Threat Protection Components

Adversaries have become highly targeted, leveraging sophisticated playbooks to breach an organization, move laterally, and extract valuable data, all while remaining invisible to traditional defenses. Threat protection automatically stops vulnerability exploits with intrusion prevention capabilities, offers in-line malware protection, and blocks outbound command-and-control-traffic. When combined with cloud-based malware analysis and Uniform Resource Locator (URL) filtering, organizations are protected at every stage of the attack lifecycle, including both known and zero-day threats.

Effective threat protection against sophisticated modern threats requires a suite of security components working together including:

- » Endpoint protection
- » Security Orchestration Automation and Response (SOAR)
- » Cloud-based threat intelligence

Endpoint protection

For decades, traditional antivirus has been the de facto solution for protecting endpoints. Antivirus checks all the boxes for regulatory, governance, and compliance audits, but it provides minimal real security benefits. Although antivirus solutions protect nearly every endpoint and server in the world, security breaches continue at an alarming rate. This is largely because traditional antivirus

is a signature-based security tool that focuses on detecting and responding to known threats after they have already entered a network.

At the same time, adversary strategies have evolved from simple malware distribution. Today, attackers can bypass antivirus with inexpensive, automated tools that produce countless unique, targeted, and sophisticated attacks. Ultimately, traditional antivirus is proving inadequate to protect systems against breaches.

Although attacks have become more sophisticated and complex, they still use basic building blocks to compromise endpoints. The primary attack methods continue to exploit known and unknown application vulnerabilities as well as deploy malicious files, including ransomware. These can be used individually or in various combinations, but they are fundamentally different in nature:

- » **Exploits** are the results of techniques used against a system that are designed to gain access through vulnerabilities in the code of an operating system or application.
- » **Malware** is a file or code that infects, explores, steals, or conducts virtually any behavior an attacker wants.
- » **Ransomware** is a form of malware that holds valuable files, data, or information for ransom, often by encrypting data, with the attacker holding the decryption key.

To effectively combat security breaches, organizations must protect themselves from known and unknown cyberthreats as well as the failures of traditional antivirus. This means they must focus on prevention — the only effective, scalable, and sustainable way to reduce the frequency and impact of cyber breaches. To deliver effective and comprehensive security to systems, endpoints, and users, endpoint protection must do the following:

- » **Preemptively block known and unknown threats.** To prevent security breaches, a shift must occur — from detecting and responding to incidents after they have occurred to preventing breaches from occurring in the first place. Endpoints must be protected from known and unknown malware and exploits, including zero-day threats, whether a machine is online or offline, on-premises or off, connected to the organization's network or not. A key step in accomplishing this is incorporating

local and cloud-based analysis to detect and prevent unknown and evasive threats.

- » **Have no negative impact on user productivity.** Advanced endpoint security must enable end users to conduct daily business as well as use mobile and cloud-based technologies without fear of unknown cyberthreats. Users should be able to focus on their responsibilities rather than worry about security patches and updates. They must be confident that they are protected from inadvertently running malware or exploits that may compromise their systems.
- » **Turn threat intelligence into prevention automatically.** Threat intelligence gained through encounters with new and unique attacks, including from third-party intelligence service providers and public intelligence-sharing constructs, must enable endpoint agents to instantly block known malware, identify and block unknown malware, and stop both from infecting endpoints. Threat data must also be gathered from the network, clouds, and endpoints within the organization. Machine learning and automation must be used to correlate the data, identify indicators of compromise, create protections, and push them out across the organization.
- » **Protect all applications.** Applications are at the core of any organization's ability to function effectively. Unfortunately, security flaws or bugs in applications create a large attack surface that traditional antivirus fails to protect. An organization's security infrastructure should be able to prevent exploitation of all third-party and proprietary applications. It should also be able to return security verdicts quickly in order to expedite approvals as new applications are introduced into the environment.
- » **Keep security out of the way of user productivity.** Breach prevention must never jeopardize user productivity, so security products should not burden computational resources. Any security, including endpoint protection, must be lightweight enough to require only minimal system resources, or it will invariably degrade the user experience and productivity.
- » **Keep legacy systems secure.** Organizations may not always deploy available system updates and security patches immediately, either because doing so would interfere with, diminish, or eliminate critical operational capabilities, or

because patches may not be available for legacy systems and software that have reached end-of-life. Therefore, a complete endpoint security solution must support systems that cannot be patched by preventing software exploits, known or unknown, regardless of the availability or application of security patches.

- » **Be enterprise ready.** Any security solution intended to replace antivirus should be scalable, flexible, and manageable enough for deployment in an enterprise environment. Endpoint security should support and integrate with the way an enterprise deploys its computing resources, scale to as many endpoints as needed, and support deployments that cover geographically dispersed environments. It must also be flexible in its ability to provide ample protection while still supporting business needs and not overly restricting the business. This flexibility is critical as the needs of one part of the organization may be entirely different from those of another. Additionally, the solution must be able to be easily managed by the same group that manages security in other parts of the organization. It must be designed with enterprise management in mind, without adding operational burden.
- » **Detect and respond to stealthy threats.** No anti-malware solution can block all endpoint threats. Adversaries, including malicious insiders, state-sponsored attackers, and advanced cybercriminals, can find underhanded ways to bypass the best malware protection. Sophisticated attackers can avoid the use of malware altogether and leverage legitimate apps and stolen credentials to execute their attacks. To find and stop attackers before the damage is done, organizations need detection powered by machine learning and integrated response to quickly contain threats. They need to identify evasive threats by continuously profiling user and endpoint behavior to uncover behavioral anomalies. The right security tool should be able to accelerate investigations by grouping related alerts into incidents and automatically revealing the root cause of any attack. Coordinated response across endpoint, network, and cloud enforcement points allows security teams to shut down threats with fast and accurate remediation.

Security orchestration and automation

Security orchestration is a method of connecting disparate security tools, teams, and infrastructures for seamless and process-based security operations and incident response. Security orchestration acts as a powerful enabler for security automation because well-connected security systems are more receptive to automation and scale.

The three pillars of security orchestration are people, processes, and technology. By streamlining security processes, connecting different security tools and technologies, and maintaining the right balance of machine-powered security automation and human intervention, security orchestration empowers security professionals to improve the organization's overall security posture.

A combination of industry trends and market forces have created challenges that security orchestration is well positioned to solve, including:

- » **Rising alert numbers:** With an increased threat surface, a greater number of entry vectors for attackers, and an increase in specialized cybersecurity tools, the number of alerts is constantly on the rise. Analysts need help in identifying false positives, duplicate incidents, and keeping the alert numbers in check without burning out.
- » **Product proliferation:** Analysts use numerous tools — both within and outside the purview of security — to coordinate and action their response to incidents. This involves lots of screen switching, fragmented information, and disjointed record keeping.
- » **Lack of skilled analysts:** With a shortage of millions of analysts expected over the coming years, many security operations centers (SOCs) are understaffed, leading to increased workload, stress, and rates of error among analysts.
- » **Inconsistent response processes:** As SOCs mature, security teams spend most of their day fighting fires and can't devote enough time to set standard response processes or spot patterns that reduce rework. This results in response quality being dependent on individual analysts, which can lead to variance in quality and effectiveness.

FIVE REQUIREMENTS FOR EFFECTIVE ENDPOINT PROTECTION

Attackers must complete a sequence of events, known as the *attack lifecycle*, to accomplish their objectives, whether stealing information or running ransomware. Nearly every attack relies on compromising an endpoint to succeed, and although most organizations have deployed some type of endpoint protection, infections are still common.

Here are five key requirements for effective endpoint protection:

- **Fighting threats with cloud-based malware analysis.** Today's complex threat landscape — combined with the diversity, volume, and sophistication of threats — makes effective threat prevention challenging. This problem is compounded by the challenge of detecting never-before-seen malware and exploits in addition to identifying known malicious content.

To address these sophisticated, targeted, and evasive threats, endpoint protection must integrate with shared, cloud-based threat intelligence to learn and evolve its defenses and enable deep analysis to rapidly detect potentially unknown threats.
- **Prevent ransomware.** Although ransomware is not new, major attacks like WannaCry, Petya/NotPetya, and TrickBot have shown that traditional prevention methods are ineffective against advanced ransomware. Attackers have evolved their approach and use of malware to become more sophisticated, automated, targeted, and highly evasive.
- **Hit pause on "Patch Tuesday."** Thousands of new software vulnerabilities and exploits are discovered each year, requiring diligent software patch distribution by software vendors on top of patch management by system and security administrators in every organization. This regular stream of patches and updates is affectionately known as "Patch Tuesday."


(continued)

(continued)

Although patching is a critical part of a sound endpoint protection strategy, it protects an organization's endpoints only after vulnerabilities are discovered and patched. Delays of days, weeks, or longer are inevitable as patches for newly discovered vulnerabilities must be developed, distributed, tested, and deployed. Much like signature-based malware detection, patch management is an endless race against time that offers no protection against zero-day exploits. Vulnerability exploits, however, are the primary reason patches are needed.

Many advanced threats work by placing malicious code (exploits) in seemingly innocuous data files. When these files are opened, the exploit leverages unpatched vulnerabilities in the native application used to view the file, and the code executes. Because the application being exploited is allowed by IT security policy, this type of attack bypasses application whitelisting controls.

Although many thousands of exploits exist, they all rely on a relatively small set of core techniques that don't frequently change. Regardless of the exploit or its complexity, for an attack to succeed, the attacker must execute a series of these core exploit techniques in sequence, like navigating a maze to reach the goal (see the figure). Thus, the key to exploit prevention is to focus on the exploit techniques, rather than the exploits themselves.



- **Protect resource-sensitive environments.** Frequent antivirus signature updates, application patches, and operating system updates required to secure endpoints against known vulnerabilities are particularly challenging in virtual environments, where “golden images” are used to provision virtual endpoints. Many traditional physical endpoint products can create unforeseen complications when applied to virtual environments. Furthermore, purpose-built virtual security products often leave gaps in the overall security architecture if they are not part of a cohesive security infrastructure.

A new approach is needed to protect virtual and cloud environments from the ground up — one that offers continuous protection without the need for signatures, patches, or updates; integrates seamlessly into any virtual environment; and is part of an end-to-end security platform that encompasses physical, virtual, and cloud-based computing environments.

- **Protect endpoints from day one.** Deploying and managing endpoint protection shouldn't be difficult. However, customers of traditional endpoint protection products complain about day-to-day management, database maintenance, agent updates, and constant tuning to eliminate false positives and keep resource utilization in check. Worse still, even with all this work, endpoints still get compromised.

Security orchestration and automation helps enterprises address these challenges with the following capabilities and benefits:

- » **Accelerate incident response:** By replacing low-level manual tasks with corresponding automations, security orchestration can shave off large chunks from incident response times while also improving accuracy and analyst satisfaction.
- » **Standardize and scale processes:** Since low-level tasks are automated and processes are standardized, analysts can spend their time in more important decision-making and charting future security improvements rather than getting mired in grunt work.
- » **Unify security infrastructures:** A security orchestration platform can act as a connective fabric that runs through previously disparate security products, providing analysts with a central console through which to action incident response.
- » **Increase analyst productivity:** Through stepwise, replicable workflows, security orchestration can help standardize incident enrichment and response processes that increases the baseline quality of response and is primed for scale.

- » **Leverage existing investments:** By automating repeatable actions and minimizing console-switching, security orchestration enables teams to coordinate among multiple products easily and extract more value out of existing security investments.
- » **Improve overall security posture:** The sum of all of these benefits is an overall improvement of the organization's security posture and a corresponding reduction in security and business risk.

Cloud-based threat intelligence

Today, organizations must contend with an entire marketplace of malware and exploit developers selling or renting out their malicious tools, making them available to all classes of attackers. At the same time, advanced evasion techniques have been commoditized, allowing attacks to sidestep legacy detection approaches. Now, even low-skilled adversaries can launch unique attacks capable of evading traditional threat identification and prevention approaches, requiring human intervention that cannot scale against the volume of unknown threats seen today.

Cloud-based threat intelligence goes beyond traditional approaches used to detect unknown threats, bringing together the benefits of four independent techniques for high-fidelity and evasion-resistant discovery, including:

- » **Dynamic analysis:** Observes files as they detonate in a purpose-built, evasion-resistant virtual environment, enabling detection of zero-day exploits and malware using hundreds of behavioral characteristics.
- » **Static analysis:** Complements dynamic analysis with effective detection of malware and exploits, as well as providing instant identification of malware variants. Static analysis further leverages dynamic unpacking to analyze threats attempting to evade detection using packer tools.
- » **Machine learning:** Extracts thousands of unique features from each file, training a predictive machine learning model to identify new malware, which is not possible with static or dynamic analysis alone.
- » **Bare metal analysis:** Detonates evasive threats in a real hardware environment, entirely removing an adversary's ability to deploy anti-VM analysis techniques.

IN THIS CHAPTER

- » Identifying applications and users
- » Leveraging threat protection and adaptive integration capabilities
- » Protecting mobile and remote users
- » Simplifying policy control and management
- » Extending security to the cloud and automating routine tasks
- » Taking advantage of flexibility in deployment options and partner integrations

Chapter 5

Ten Evaluation Criteria for Network Security in the Data Center and Hybrid Cloud Environment

This chapter helps you assess network security solutions for your data center or hybrid cloud environment by presenting several important features and criteria for you to consider.

Safe Enablement of Applications in the Hybrid Cloud

More and more applications, such as instant messaging (IM) applications, peer-to-peer (P2P) file sharing, or Voice over IP (VoIP), are capable of operating on nonstandard ports or hopping ports. Additionally, users are accessing diverse types of apps, including software-as-a-service (SaaS) apps, from varying devices and locations. Some of these apps are sanctioned, some tolerated, and others unsanctioned, and users are increasingly savvy enough to force applications to run over nonstandard ports through protocols such as Remote Desktop Protocol (RDP) and Secure Shell (SSH).

Furthermore, new applications provide users with rich sets of functions that help ensure user loyalty but may represent different risk profiles. For example, Webex is a valuable business tool, but using Webex desktop sharing to take over an employee's desktop from an external source may be an internal or regulatory compliance violation. Gmail and Google Drive are other good examples. Once users sign in to Gmail, which may be allowed by policy, they can easily switch to YouTube or Google Photos, which may not be allowed.

Another common practice in hybrid environments is to mix application workload trust levels on the same compute resources. Although efficient in practice, mixed levels of trust introduce additional security risks in the event of a compromise. Your network security platform must be able to implement security policies based on the concept of Zero Trust as a means of controlling traffic between workloads (segmentation of east-west traffic) while preventing lateral movement of threats.

Security administrators need to have complete control over usage of these apps and must be able to set policy to allow or control certain types of applications and application functions while denying others. Your network security platform for the data center and hybrid cloud environment must classify traffic by application on all ports, all the time, by default — and it should not burden you with researching common ports used by each application. It must provide complete visibility into application usage along with capabilities to understand and control their use (see Figure 5-1).




FIGURE 5-1: Control application usage in policy.

For example, your network security platform should understand usage of application functions, such as audio streaming, remote access, and posting documents, and be able to enforce granular controls over that usage, such as upload versus download permissions, chat versus file transfer, and so on. This must be done continuously. The concept of “one-and-done” traffic classification is not an option because it ignores the fact that these commonly used applications share sessions and support multiple functions. If a different function or feature is introduced in the session, such as sharing a desktop in a Webex conference, the network security platform must perform a policy check again. Continuous state tracking to understand the functions each application may support — and the different associated risks — is a must.

Once a complete picture of applications is gained, safe application enablement of applications is essential to deliver the right security policies in the data center or across a hybrid cloud environment. This includes more fine-grained and appropriate application functions than simply “allow” or “deny,” such as allow but enforce traffic shaping through Quality of Service (QoS) or allow based on schedule, users, or groups. Application visibility and control allows organizations to reduce the attack surface by blocking rogue and misconfigured applications, such as unauthorized management tools and P2P file-sharing software. It also enables the protection of high-value targets, such as domain controllers, finance servers, and email and database servers with meaningful network segmentation.



REMEMBER

Accurate traffic classification — regardless of ports, protocols, evasive tactics, and Secure Sockets Layer (SSL) encryption — is important in any data center. This is even more critical in a hybrid cloud environment where virtual machines (VMs) and application workloads communicate between on-premises and cloud environments, often without appropriate policies or risk analysis.

Machine learning is helpful for accurate traffic classification due to its ability to continuously adapt to data patterns and environment variables. However, this level of classification is challenging, and many organizations take a more general, phased approach with less fine-grained micro-segmentation.

Identify Users and Enable Appropriate Access

Employees, customers, and partners connect to different repositories of information within your network, as well as to the Internet. These people and their many devices represent your network's users. It's important for your organization's risk posture that you're able to identify your users beyond Internet Protocol (IP) address, as well as grasp the inherent risks they bring, based on the devices they're using — especially when security policies have been circumvented or new threats have been introduced to your network. In addition, users are constantly moving to different physical locations and using multiple devices, operating systems, and application versions to access the data they need (see Figure 5-2). IP address subnets are mapped only to physical devices, not individual users, meaning that if users move around — even within the office — policy doesn't follow them. Therefore, user and group information must be directly integrated into the technology platforms that secure data centers and hybrid cloud environments.

Your network security platform must be able to pull user identity from multiple sources, including virtual private networks (VPNs), wireless local area network (WLAN) access controllers, directory servers, email servers, and captive portals. Knowing who is using the applications on your network, and who may be transmitting a threat or transferring files, strengthens security policies, and improves incident response times. The platform must allow policies to safely enable applications based on users or groups of users, outbound or inbound — for example, by allowing only your IT department to use tools such as SSH, Telnet, and File Transfer Protocol (FTP). User-based policies follow users no matter where they go — at headquarters, branch offices, or home — and on whatever devices they use. However, the issue of user identity goes beyond classifying users for policy reporting.

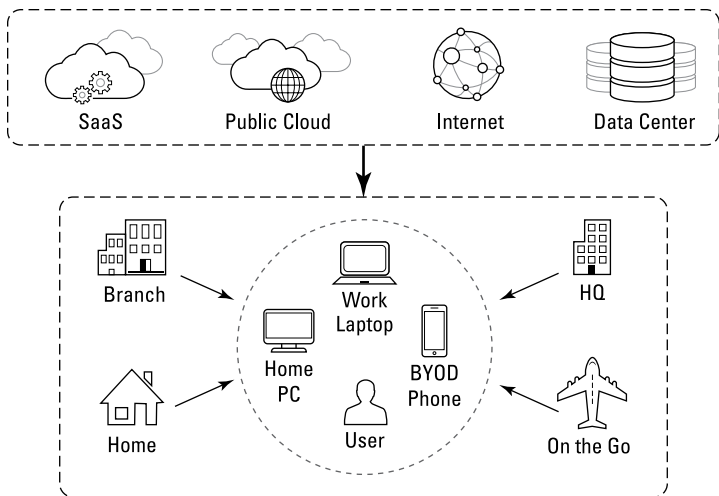


FIGURE 5-2: Users access data from different devices and locations.

Comprehensive Threat Protection

The modern threat landscape has evolved into intelligent, targeted, persistent, multiphase intrusions. Threats are delivered via applications that dynamically hop ports, use nonstandard ports, tunnel within other applications, and hide within proxies, SSL, or other types of encryption. Within the data center, exerting application-level control between your workloads reduces your threat footprint while simultaneously segmenting data center traffic based on Zero Trust principles. Application-specific threat prevention policies can prevent known and unknown threats from compromising your data center.

Additionally, enterprises are exposed to targeted and customized malware, which can easily pass undetected through traditional port-based firewalls and antivirus software. Most modern malware — including ransomware variants — uses advanced techniques, such as wrapping malicious payloads in legitimate files or packing files to avoid detection, to transport attacks or exploits through network security devices and tools. As organizations have increasingly deployed virtual sandboxes for dynamic analysis, attackers have evolved to focus on ways to evade them. They employ techniques that scan for valid user activity, system configurations, or indicators of specific virtualization technologies. With the growth of the cybercrime underground, any attacker, novice or

advanced, can purchase plug-and-play threats designed to identify and avoid malware analysis environments.

In addition, file and data filtering options — for example, the ability to block files by their actual type and the ability to control the transfer of sensitive data patterns, such as credit card numbers — address important compliance use cases.

One of the limitations of traditional antimalware security signatures is the ability to protect only against malware that has been previously detected and analyzed. This reactive approach creates a window of opportunity for malware. To supplement this, the data center network security solution should provide the ability to directly analyze unknown executables for malicious behavior.

Your network security platform, using integrated security services, should automatically block known threats. Unknown threats must be automatically analyzed and countered, too. Your organization needs a service that looks for threats at all points within the cyberattack life cycle (see Figure 5-3), not just when threats first enter your network. Blocking known risky file types or access to malicious Uniform Resource Locators (URLs) before they compromise your network reduces your threat exposure. Your network security platform should protect you from known vulnerability exploits, malware, and command-and-control (C2) activity without requiring you to manage or maintain multiple single-function appliances. Signatures should be updated automatically as soon as new malware is encountered, keeping you protected while allowing your security and incident response teams to focus on the things that matter.




FIGURE 5-3: Disruption at every step to prevent successful attacks.

A network security platform that utilizes multiple methods of analysis to detect unknown threats, including static analysis with machine learning, dynamic analysis, and bare metal analysis, is capable of high-fidelity, evasion-resistant discovery. Rather than use signatures based on specific attributes, your network security

platform should use content-based signatures to detect variants, polymorphic malware, or C2 activity. In addition, C2 signatures based on analysis of outbound communication patterns are much more effective protective measures that can scale at machine speed when created automatically. Finally, cloud-delivered security infrastructure is critical for security enforcement. It supports threat detection and prevention at massive scale across your network, endpoints, and clouds in addition to allowing you to tap into an open ecosystem of trusted innovators.

Flexible, Adaptive Integration

One of the key integration challenges in the data center is security design. Network architectures must often be redesigned when security requirements evolve due to changing applications and threats, new compliance mandates, and shifting risk postures. A new paradigm that enables network security to be flexible and adaptive is needed.

Networking flexibility helps ensure compatibility with practically any organization's data center environment. Enabling integration without the need for redesign or reconfiguration depends not only on supporting a wide range of networking features and options, such as port-based virtual local area networks (VLANs), but also on the ability to integrate at the Open Systems Interconnection (OSI) Reference Model Layer 1 (Physical), Layer 2 (Data Link), or Layer 3 (Network). In addition, the network security solution should be able to turn on additional security features as the security posture changes. Finally, your security solution needs the ability to support multiple hypervisor types, such as VMware ESXi, Microsoft Hyper-V, Nutanix AHV, KVM, and potentially software-defined network (SDN) driven provisioning, particularly in hybrid cloud environments.

Secure Access for Mobile and Remote Users

The modern enterprise continues to become far more distributed than in the past. The mobile workforce continues to grow along with the use of mobile devices to connect to business applications,

often through public networks and devices that are open to advanced threats. Users simply expect to be able to connect and work from any location, whether at an airport, in a coffee shop, in a hotel room, or at home. Employees, partners, contractors, and supply chains are all accessing data center and cloud resources from beyond the traditional perimeter of the enterprise. This increases risk when users are off-premises because there is no network firewall to stop attacks. The issue becomes even more complex when considering the effects of cloud and bring-your-own-device (BYOD) practices. In addition, remote locations and small branch offices often lack consistent security because it is operationally inefficient and costly to ship firewalls to them or backhaul traffic to headquarters.

The mobile workforce and remote locations need access to applications from places far beyond your network. The requirement to protect these mobile and remote users is a way to enable the same application, user, and content protections they receive while on premises. They also need protection from targeted cyberattacks, malicious applications and websites, phishing, C2 traffic, and other unknown threats. This requires consistent security. Your network security platform must enable and adapt to the required levels of visibility, threat prevention, and security policy enforcement required to protect distributed users and locations by delivering network security capabilities from the cloud across physical and virtualized workloads spanning hybrid cloud, securing them with either physical hardware or virtualized appliances.



REMEMBER

Network security solutions for the hybrid data center must deliver secure access for mobile and remote users to the data center and cloud, in addition to addressing the use of endpoint devices other than standard corporate-issued equipment.

One Comprehensive Policy, One Management Platform

Individual security products typically come with their own management applications. To configure security for each one, security operators must work with different management systems. These products are often disconnected and cannot share insights. Organizations also find it challenging to scale firewall onboarding, maintain consistent security policies, and deploy emergency changes across thousands of firewalls. This makes security complex and stretches IT teams to the limit.

Organizations must be able to operationalize the deployment of consistent, centralized security policies across tens of thousands of firewalls spanning on-premises and cloud deployments — including remote locations, mobile users, and SaaS applications — through centralized management, consolidated core security tasks, and streamlined capabilities. For example, you should be able to use a single console to view all network traffic, manage configuration, push global policies, and generate reports on traffic patterns or security incidents. Your reporting capabilities must let your security personnel rapidly drill down into network, application, and user behavior for the context they need to make informed decisions.

When these capabilities are delivered from the cloud, your teams can build out the right security architecture to prevent known and unknown threats at every corner of your extended network. In today's constantly changing threat landscape, using a single security vendor to address the vast spectrum of your security and business needs isn't always practical. In this case, the ability to integrate with and consume third-party and cloud service provider (CSP) insight and telemetry is critical. When evaluating future security vendors, be sure to evaluate the integration, extensibility, and programmability that they offer.



REMEMBER

Hybrid data centers are composed of physical and virtual infrastructures deployed on-premises and in private, public, and hybrid cloud environments. Network security solutions in the hybrid data center need to include both physical and virtualized options. The network security policy management platform must also support hybrid data center environments; otherwise, security policies can become convoluted, leading to needless complexity, misconfigurations, and security blind spots. In addition, a single, comprehensive security policy that fully integrates application control, threat management, and user identification is a must.

Cloud Ready

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, and minimizes the security policy lag as virtual workloads change. Your network security platform 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, it must support a variety of cloud and virtualization environments, including all major public cloud providers and virtualized private clouds. The network security platform must integrate with native cloud services — such as Amazon Lambda and Azure Kubernetes Service, Azure Functions, and Azure App Service — as well as automation tools, such as Ansible and Terraform, to integrate security into your cloud-first development projects.

Data center tasks and processes that help IT teams execute change with greater speed, quality, and consistency are typically automated using workflows. However, deployment of security capabilities typically lags orchestration software provisioning in virtual and cloud environments, leading to security risks and considerable integration challenges. Automated provisioning of network security capabilities, in line with other orchestration elements of the hybrid data center environment, is essential.

Automate Routine Tasks and Focus on the Threats That Matter

A recent survey from the Enterprise Strategy Group (ESG) found 51 percent of cybersecurity professionals feel their organization has a problematic shortage of cybersecurity skills. This is compounded by a dependency on too many manual processes for day-to-day security operations, such as chasing down data, investigating false positive alerts, and managing remediation. Manually analyzing and correlating the vast number of security events slows mitigation, increases the chance for error, and is difficult to scale. Security teams can easily drown in the volume of alerts and miss the critical, actionable ones. This is exacerbated by a looming shortage of skilled cybersecurity professionals. Although big data analytics uncovers hidden patterns, correlations, and other insights to provide security teams with actionable intelligence, you still need the right data. That data must be sourced from everywhere — networks, endpoints, SaaS applications, public clouds, private clouds, data centers, and so on — and be ready for analytics.

By using precise analytics to drive automation, you can easily operate security best practices like Zero Trust; streamline routine tasks; and focus on business priorities, such as speeding

application delivery, improving processes, or hunting for threats. There are three ways to think about automation:

- » **Workflow automation.** Your network security platform must expose standard application programming interfaces (APIs) so it can be programmed from other tools and scripts you may be using. In the cloud, it must integrate with tools like Ansible and Terraform. In addition, it must be able to kick off workflows on other devices in your security ecosystem, using their APIs, without manual intervention.
- » **Policy automation.** The network security platform must be able to adapt policies to any changes in your environment, such as movement of applications across virtual machines. It must also be able to ingest threat intelligence from third-party sources and automatically act on that intelligence.
- » **Security automation.** Your environment must be able to uncover unknown threats and deliver protections to the network security platform so new threats are blocked automatically.

Some threats remain hidden in data. By looking deeper into that data across locations and deployment types, you can find threats that may be lurking in plain sight. With automation, you can accurately identify threats, enable rapid prevention, improve efficiency, better utilize the talent of your specialized staff, and improve your organization's security posture.

Flexible Deployment Options

The choice of whether a physical or virtual network security appliance should be deployed in the data center depends on the specific issues to be addressed.

Physical network security appliances are often adequate if the same trust levels are maintained within a single cluster of virtual hosts. In this scenario, visibility of east-west traffic (internal communications between servers) is less critical and can be forced off-box through a default security appliance, if necessary. Virtual systems also offer scale and performance in the data center with east-west traffic inspection, and larger physical platforms can be partitioned with virtual systems for both north-south and east-west traffic inspection.

However, when applications of different trust levels exist within the same virtual cluster (for example, VMware and Nutanix), full visibility of intra-host communications can be achieved only with virtual firewalls.

Additionally, specific server-level or hypervisor attacks can only be addressed with firewalls protecting these servers — either virtual or physical.

Finally, firewalls deployed to the public cloud may need to be virtualized if the responsible service providers don't allow customers to deploy their physical hardware in the cloud. Virtualized firewalls may also be appropriate in private clouds and hybrid data centers, where rack space is at a premium or where data center mobility (for example, a temporary data center in a remote region) is needed.

Consume New Innovations Easily in a Broad Partner Ecosystem

Consuming cybersecurity innovation is arduous. 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.

As the number of needed security functions increases, there are two options: Add more independent point products or use an existing device to support new capabilities. If your network security platform can act as a sensor and enforcement point for third-party technology, you can rapidly adopt new security innovations without deploying or managing endless new devices. Your network security platform 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, and drive automated response and enforcement with deeply integrated applications. This way, they can solve the most challenging security use cases with the best technology available, and they can do so without the cost or operational burden of deploying new infrastructure for each new function.

Glossary

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

advanced persistent threat (APT): A sustained Internet-borne attack, usually perpetrated by a group with significant resources, such as organized crime or a nation-state.

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

Australian Privacy Principles: The Privacy Act 1988 establishes standards for collecting and handling personal information, referred to as the Australian Privacy Principles (APP).

blended threat: Transforming a threat by making a relatively minor change to malicious code or by adding entirely new propagation and exploit mechanisms.

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

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

brute force: A type of attack in which the attacker attempts every possible combination of letters, numbers, and characters to crack a password, passphrase, or PIN.

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

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.

continuous delivery: Software code in a CI pipeline must go through manual technical checks before it is implemented in production. *See also* continuous integration (CI).

continuous deployment (CD): Software code in a CI pipeline passes automated testing rather than manual checks (as in continuous delivery) is automatically deployed, giving customers instant access to new features. *See also* continuous delivery *and* continuous integration (CI).

continuous integration (CI): An automated software development pipeline which requires developers to integrate code into a repository frequently (for example, several times per day) for automated testing. Each check-in is verified by an automated build, allowing teams to detect problems early.

Control Panel File (CPL): The filename extension for Control Panel items in Microsoft Windows.

cryptocurrency: A form of digital currency, such as Bitcoin, that uses encryption to control the creation of currency and verify the transfer of funds independent of a central bank or authority.

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.

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

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

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

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

dynamic address group: An object group, used in policies, that uses tags as filtering criteria to determine its members. A dynamic address group allows you to create policies that adapt to changes such as adds, moves, or deletions of VMs. *See also* virtual machine (VM).

Dynamic Host Configuration Protocol (DHCP): A standardized protocol that provides TCP/IP and Link Layer configuration parameters and network addresses to dynamically configured hosts on a TCP/IP network. *See also* Transmission Control Protocol (TCP) *and* Internet Protocol (IP).

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

file transfer protocol (FTP): A standardized protocol used for the transfer of files between a client and server over a network.

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

Health Insurance Portability and Accountability Act (HIPAA): U.S. legislation passed in 1996 that, among other things, protects the confidentiality and privacy of protected health information (PHI). *See also* protected health information.

hybrid cloud: An environment 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.

HyperText Transfer Protocol (HTTP): The primary communication protocol of the Internet.

HyperText Transfer Protocol Secure (HTTPS): A secure communication protocol widely used on the Internet.

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

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

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

intellectual property (IP): Proprietary information including patents, trademarks, copyrights, and trade secrets.

Internet of Things (IoT): The network of physical smart, connected objects that are embedded with electronics, software, sensors, and network connectivity.

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

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

Lightweight Directory Access Protocol (LDAP): An IP and data storage model that supports authentication and directory functions. *See also* Internet Protocol (IP).

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

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

metamorphism: A technique used in a virus to change its appearance in host programs without necessarily depending on encryption. The difference in appearance comes from changes made by the virus to its own body.

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). 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.

multi-factor authentication (MFA): Any authentication mechanism that requires two or more of the following factors: something you know, something you have, something you are.

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

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

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

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

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

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

point-of-sale (POS) systems: A system, such as a cash register or credit card terminal, used to complete a sales transaction.

polymorphism: A technique used in a virus to change its appearance in host programs. For instance, it encrypts its body with a different key each time and prepends a decryption routine to itself. The decryption routine (known as the “decryptor”) is mutated randomly across virus instances, so as to be not easily recognizable.

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

protect surface: The critical data, applications, assets, and services that need to be protected in a Zero Trust architecture. *See also* Zero Trust.

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

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

Quality of Service (QoS): The ability to prioritize various types of voice and data traffic based on operational needs such as response time, packet loss, and jitter.

ransomware: Malware that encrypts files on an infected server or endpoint and demands a ransom payment, usually cryptocurrency, to retrieve the key to decrypt the files.

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

Representational State Transfer (REST): A software architectural style that defines a set of constraints to be used for creating web services.

rootkit: Malware that provides privileged (root-level) access to a computer. *See also* malware.

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

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

Server Message Block (SMB): An application-layer protocol also known as Common Internet File System (CIFS).

social engineering: A technique commonly used in phishing attacks that relies on implicit trust to extract sensitive information from a victim.

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

software-defined networking (SDN): An approach to networking that uses virtualization to abstract higher-level network services from underlying physical hardware.

spear phishing: A phishing attack that's highly targeted; for example, at a particular organization or part of an organization. *See also* phishing.

supervisory control and data acquisition (SCADA): An industrial automation system that operates with coded signals over communication channels to provide remote control of equipment.

telnet: A network protocol used to establish a command line interface on another system over a network.

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

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

virtual local area network (VLAN): A LAN segment that is partitioned by broadcast domain at Layer 2 (Data Link) of the OSI model, typically configured on a switch or router. *See also* local area network (LAN) and Open Systems Interconnection (OSI) Reference Model.

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

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

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

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

web application firewall (WAF): A device used to protect a web server from web application attacks such as script injection and buffer overflow.

wireless local area network (WLAN): A Wi-Fi network in a relatively limited geographical area such as an office building, coffee shop, hotel, or airport.

worm: A type of malware that spreads copies of itself from computer to computer.

ws2_32.dll: A dynamically linked library that is used to handle network connections.

Zero Trust: A strategic initiative, rooted in the principle of “never trust, always verify,” that helps prevent data breaches by eliminating inherent trust on your network and in your data center in order to restrict unauthorized lateral movement.

You deserve stronger, simpler security.

Data centers are becoming more complex. What's more, security must follow workloads everywhere – on-prem, in the cloud and everywhere in between.

By partnering with Palo Alto Networks for data center security and cloud transformation, you can overcome these challenges. Find out how at paloaltonetworks.com/network-security/data-center.

Take a virtual Ultimate Test Drive and experience our best-of-breed security products firsthand from your home or office. Visit paloaltonetworks.com/events/test-drive for details.



Safeguard your data center and hybrid cloud

The evolution of the data center and hybrid cloud helps IT organizations deliver greater business opportunities, but also introduces new risks. Data centers and hybrid clouds that span multi-cloud environments offer a larger attack surface, which can translate to increased complexity in networking and cybersecurity. It is critical to maintain full visibility and precise control of your data center regardless of the architecture. This book is your guide to protecting your data center and hybrid cloud environments.

Inside...

- Gain complete visibility
- Minimize the attack surface
- Automate threat protection
- Implement a shared responsibility model
- Protect endpoints
- Deploy cloud-based threat intelligence
- Help prevent data breaches with a Zero Trust security plan



Lawrence C. Miller has worked in information technology for more than 25 years. He has written almost 200 For Dummies books.

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

ISBN: 978-1-119-68155-7
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.