

**Making Everything Easier!™**

**Palo Alto Networks® 2nd Edition**

# **Network Security in Virtualized Data Centers**

FOR  
**DUMMIES®**  
A Wiley Brand

## **Learn to:**

- Securely enable applications in public, private, and hybrid cloud environments
- Ensure visibility and control of east-west traffic in the data center
- Implement a phased approach to virtualized data center security

*Brought to you by*



**Lawrence C. Miller, CISSP**



# Palo Alto Networks

Palo Alto Networks is leading a new era in cybersecurity by securing thousands of enterprise, government, and service provider networks from cyber threats and protecting our digital way of life. The next-generation platform uses an innovative traffic classification engine that identifies network traffic by application, user, and content.

The Palo Alto Networks next-generation security platform is built on four main principles:

1. **Natively integrated** technologies that support open communication, orchestration, and visibility;
2. **Automation** of protection creation and reprogramming of the security posture across network, endpoint and cloud environments;
3. **Extensibility** that allows for protection of customers as they expand and market requirements change; and
4. **Threat intelligence sharing** to minimize the spread of attacks by providing protection based on comprehensive global threat data.

The next generation security platform offers superior protection against the sophistication of modern attacks, can reduce the total cost of ownership for organizations by simplifying their security infrastructure, and eliminates the need for multiple, stand-alone security appliances and software products.

Find out more at [www.paloaltonetworks.com](http://www.paloaltonetworks.com)

# ***Network Security in Virtualized Data Centers***

FOR  
**DUMMIES**<sup>®</sup>  
A Wiley Brand

***Palo Alto Networks<sup>®</sup> 2nd Edition***



# ***Network Security in Virtualized Data Centers***

FOR  
**DUMMIES**<sup>®</sup>  
A Wiley Brand

***Palo Alto Networks<sup>®</sup> 2nd Edition***

**by Lawrence C. Miller, CISSP**

FOR  
**DUMMIES**<sup>®</sup>  
A Wiley Brand

## Network Security in Virtualized Data Centers For Dummies®, Palo Alto Networks® 2nd Edition

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

Copyright © 2016 by John Wiley & Sons, Inc., Hoboken, New Jersey

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

**Trademarks:** Wiley, For Dummies, the Dummies Man logo, The Dummies Way, Dummies.com, Making Everything Easier, and related trade dress are trademarks or registered trademarks of John Wiley & Sons, Inc. and/or its affiliates in the United States and other countries, and may not be used without written permission. All other trademarks are the property of their respective owners. John Wiley & Sons, Inc., is not associated with any product or vendor mentioned in this book.

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

ISBN 978-1-119-24978-8 (pbk); ISBN 978-1-119-24979-5 (ebk)

Manufactured in the United States of America

10 9 8 7 6 5 4 3 2 1

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](mailto:info@dummies.biz), or visit [www.wiley.com/go/custompub](http://www.wiley.com/go/custompub). For information about licensing the *For Dummies* brand for products or services, contact [BrandedRights&Licenses@Wiley.com](mailto:BrandedRights&Licenses@Wiley.com).

## Publisher's Acknowledgments

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

**Development Editor:** Elizabeth Kuball

**Copy Editor:** Elizabeth Kuball

**Acquisitions Editor:** Amy Fandrei

**Editorial Manager:** Rev Mengle

**Business Development Representative:**  
Karen Hattan

**Production Editor:** Antony Sami

# Table of Contents

<b>Introduction</b> .....	<b>1</b>
About This Book .....	1
Foolish Assumptions .....	1
Icons Used in This Book.....	2
Beyond the Book.....	2
Where to Go from Here .....	3
<b>Chapter 1: Data Center Evolution</b> .....	<b>5</b>
The Journey to the Cloud .....	5
Why Is the Data Center Evolving?.....	8
<b>Chapter 2: Security Challenges in the Cloud</b> .....	<b>11</b>
Security Challenges with Applications in the Data Center .....	12
Applications Aren't All Good or All Bad.....	13
Hiding in Plain Sight.....	14
Threats Are Taking a Free Ride.....	17
Management applications.....	18
Unknown applications.....	19
Recognizing Data Center Challenges .....	20
Data Center Attack Vectors .....	21
Recognizing Key Characteristics of Advanced Malware.....	22
Malware Threats to the Data Center .....	24
Targeted intrusions .....	24
Advanced persistent threats.....	25
Know Thy Enemy .....	26
Understanding Modern Cyberattack Strategy.....	28
<b>Chapter 3: Securing Public and Private Clouds</b> .....	<b>33</b>
Public and Private Cloud Security Challenges .....	33
Securing the Cloud with Next-Generation Firewalls.....	36
Ensuring visibility and control.....	37
Preventing unknown traffic .....	37
Protecting users anywhere.....	38
Controlling east-west traffic .....	39
Keeping pace with the business .....	42



**Chapter 4: Securing the Hybrid Cloud . . . . .45**

- A Phased Approach to Security in Virtualized Data Centers ..... 45
  - Phase 1: Consolidating servers within trust levels ..... 46
  - Phase 2: Consolidating servers across trust levels ..... 46
  - Phase 3: Selective network security virtualization.... 47
  - Phase 4: Dynamic computing fabric..... 48
- Journey to the Cloud — One Step at a Time ..... 48
- Hybrid Cloud Use Cases ..... 50
  - Development and test ..... 51
  - Capacity augmentation ..... 51
  - Limiting “shadow” IT ..... 52

**Chapter 5: Ten (Or So) Evaluation Criteria for Network Security in the Virtualized Data Center . . . . .55**

- Safe Application Enablement of Data Center Applications ..... 55
- Identification Based on Users, Not IP Addresses..... 56
- Comprehensive Threat Protection ..... 56
- Flexible, Adaptive Integration ..... 57
- Secure Access for Mobile and Remote Users..... 58
- One Comprehensive Policy, One Management Platform.... 58
- Cloud-Readiness..... 59
- Flexible Deployment Options ..... 59

**Glossary . . . . .61**



# Introduction

---

**V**irtualization in public, private, and hybrid cloud environments has become a powerful engine for driving business growth, enabling business agility, and drastically reducing time to market for new applications and services. Organizations today are moving beyond basic server and workload consolidation initiatives and fully leveraging virtualization technologies in their cloud strategies to build a real competitive advantage.

Yet the very benefits of virtualization — for example, the ability to provide self-service resource activation and improve IT responsiveness to business demands — also introduce a myriad of security complexities. These include having visibility into virtual machine (VM) traffic that may not leave the virtual infrastructure, the ability to tie security policies to VM instantiation and movement, and segmentation of virtual machines with different trust levels.

## *About This Book*

Tackling the security implications for a virtualized computing environment is essential for the journey to the public, private, or hybrid cloud. This book outlines the challenges of securing the virtualized data center and cloud computing environments and how to address them with next-generation firewalls.

Virtualization topics cover many technologies, including servers, storage, desktops, and applications, among others. The focus of this book is network security in the virtualized data center — specifically, server virtualization.

## *Foolish Assumptions*

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

I assume that you know a little something about server virtualization, network security, and firewalls. As such, this book is written primarily for technical readers who are evaluating network security solutions to address modern threats and challenges in virtualized data centers.

## Icons Used in This Book

Throughout this book, you occasionally see special icons that call attention to important information. You won't find smiley faces winking at you or any other cute little emoticons, but you'll definitely want to take note! Here's what you can expect:



This icon points out information that may well be worth committing to your nonvolatile memory, your gray matter, or your noggin — along with anniversaries and birthdays!



You won't discover a map of the human genome here (or maybe you will, hmm), but 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 legends — well, nerds — are made of!



Thank you for reading, hope you enjoy the book, please take care of your writers! Seriously, this icon points out helpful suggestions and useful nuggets of information.



This is the stuff your mother warned you about . . . well, okay — probably not. But these helpful alerts do offer practical advice to help you avoid making potentially costly mistakes.

## Beyond the Book

There's only so much I can cover in 72 short pages, so if you find yourself at the end of this book, thinking, "Gosh, this was an amazing book, where can I learn more?" just go to [www.paloaltonetworks.com](http://www.paloaltonetworks.com).

## *Where to Go from Here*

With our apologies to Lewis Carroll, Alice, and the Cheshire cat:

“Would you tell me, please, which way I ought to go from here?”

“That depends a good deal on where you want to get to,” said the Cat — er, the Dummies Man.

“I don’t much care where . . . ,” said Alice.

“Then it doesn’t matter which way you go!”

That’s certainly true of *Network Security in the Data Center For Dummies*, which, like *Alice in Wonderland*, is also destined to become a timeless classic!

If you don’t know where you’re going, any chapter will get you there — but Chapter 1 is a good place to start! However, if you see a particular topic that piques your interest, feel free to jump ahead to that chapter. Each chapter is individually wrapped (but not packaged for individual sale) and written to stand on its own, so feel free to start reading anywhere and skip around! Read this book in any order that suits you (though I don’t recommend upside down or backward). I promise you won’t get lost falling down the rabbit hole!

# 4

## Network Security in Virtualized Data Centers For Dummies \_\_\_\_\_

# Chapter 1

---

# Data Center Evolution

.....

## *In This Chapter*

- ▶ Transforming the physical data center
  - ▶ Addressing business challenges with virtualization and cloud technologies
  - ▶ Protecting data in the cloud
- .....

**T**his chapter explains how virtualization has transformed the modern data center and ushered in the cloud era. To understand why this transformation is happening, I examine some of the benefits of virtualization and cloud technologies. Finally, I briefly discuss data security in the cloud before delving much deeper into cloud security in Chapter 2.

## *The Journey to the Cloud*

Virtualization enables organizations to utilize data center infrastructure more effectively, which helps lower costs and improves operational efficiencies. Virtualization initiatives often begin by consolidating applications on underutilized server hardware within an enterprise data center. These efforts often expand beyond consolidation to include virtualization of storage, networking, and other physical infrastructure, in order to realize other virtualization benefits, and the enterprise data center thus becomes a private cloud.

As enterprise IT needs continue to evolve toward on-demand services, many organizations move beyond data center virtualization and the private cloud, and embrace public cloud-based services and infrastructure. Popular public cloud solutions include Software as a Service (SaaS), such as Salesforce, and “prebuilt, on-demand” Infrastructure as a Service (IaaS), such as Amazon Web Services (AWS).

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 cloud computing.

## A few definitions to make the cloud a little less “cloudy”

It seems that practically everyone is talking about the cloud today, but the cloud is many things to many people. So to “clear the air,” let's start with a few standard definitions, and who better to define standards than the U.S. National Institute of Standards and Technology (NIST)?

NIST defines cloud computing as

a model for enabling ubiquitous, convenient, on-demand network access to a shared pool of configurable computing resources (such as networks, servers, storage, applications, and services) that can be rapidly provisioned and released with minimal management effort or service provider interaction.

The NIST cloud model is composed of five essential characteristics, three service models, and four deployment models.

The five essential characteristics of cloud computing are

✔ **On-demand self-service:** Computing capabilities (such as server resources) can be

unilaterally and automatically provisioned without service provider human interaction.

✔ **Broad network access:** Services are available over the network through various platforms, such as PCs, laptops, smartphones, and tablets.

✔ **Resource pooling:** Computing resources (such as processing, memory, storage, and network bandwidth) are dynamically assigned and reassigned according to demand and pooled to serve various customers (multitenancy).

✔ **Rapid elasticity:** Capabilities can be provisioned and released, in some cases automatically, to scale with demand.

✔ **Measured service:** Resource usage can be monitored, controlled, optimized, and reported.

Virtualization is the fundamental cloud-enabling technology that delivers the five essential characteristics of the cloud computing model in the virtualized data center.

The three service models defined for cloud computing include

- ✔ **Software as a Service (SaaS):** Customers are provided access to an application running on a cloud infrastructure. The application is accessible from various client devices and interfaces, but the customer has no knowledge of, and does not manage or control, the underlying cloud infrastructure. The customer may have access to limited user-specific application settings.
- ✔ **Platform as a Service (PaaS):** Customers can deploy supported applications onto the provider's cloud infrastructure, but the customer has no knowledge of, and does not manage or control, the underlying cloud infrastructure. The customer has control over the deployed applications and limited configuration settings for the application-hosting environment.
- ✔ **Infrastructure as a Service (IaaS):** Customers can provision processing, storage, networks, and other computing resources and deploy and run operating systems and applications, but the customer has no knowledge of, and does not manage or control, the underlying cloud infrastructure. The customer has control over operating systems, storage, and deployed applications, as well as some

networking components (for example, host firewalls).

Finally, NIST defines four cloud computing deployment models:

- ✔ **Public:** A cloud infrastructure that is open to use by the general public. It's owned, managed, and operated by a third party (or parties) and exists on the cloud provider's premises.
- ✔ **Community:** A cloud infrastructure that is used exclusively by a specific group of organizations. The community cloud model isn't as common as the other cloud models, so I don't specifically address it throughout this book. However, all the topics covered in this book are applicable to community cloud models as well.
- ✔ **Private:** A cloud infrastructure that is used exclusively by a single organization. It may be owned, managed, and operated by the organization or a third party (or a combination of both), and may exist on or off premises.
- ✔ **Hybrid:** A cloud infrastructure that is composed of two or more of the aforementioned deployment models, bound together by standardized or proprietary technology that enables data and application portability (for example, failover to a secondary data center for disaster recovery or content delivery networks across multiple clouds).

# Why Is the Data Center Evolving?

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 both 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.

The main driver for moving to a cloud computing strategy in the data center today is business agility. 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:

- ✔ **Improves performance:** Many applications experience asynchronous or bursty demand loads. Virtualization technologies, such as resource schedulers and VM migrations, provide intelligent, automated management of such applications. This prevents resource 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. 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.
- ✔ **Lowers costs:** IaaS and PaaS offerings in public cloud environments, and virtualization in a private cloud, enable organizations to reduce capital investments for new infrastructure, as well as operating expenses such as power, cooling, and rack space in the data center.

# **10** Network Security in Virtualized Data Centers For Dummies \_\_\_\_\_

## Chapter 2

# Security Challenges in the Cloud

---

### *In This Chapter*

- ▶ Identifying applications as good, bad, or good and bad
  - ▶ Understanding accessibility and evasion tactics
  - ▶ Recognizing how threats use applications as an attack vector
- 

**T**he purpose of a data center is to serve up applications. In the data center, business applications constitute good traffic that should be allowed on the network; other nonbusiness applications potentially constitute bad traffic that should be blocked from the network.

However, the lines between business applications and nonbusiness applications are blurring. Beyond core business applications in the data center, many cloud-based, business-supporting applications are being used in organizations. For example, storage applications such as Box, Dropbox, and OneDrive, and backup applications such as Carbonite, CrashPlan, and Mozy are now quite common. The ability to classify types of applications as good or bad is no longer a relatively straightforward exercise. Understanding the changes in the application landscape and the threat vector that these applications bring is essential to understanding what types of enablement policies are appropriate for the data center. Instead of arbitrarily blocking such applications, the application connectivity needs to be curated and controlled.

This chapter explores the security challenges with application enablement in the data center and dives into the new application landscape.

# *Security Challenges with Applications in the Data Center*

Application developers in the data center are challenged with delivering and supporting hundreds, if not thousands, of applications to employees. As business needs evolve, applications continue to be developed and improved to meet specific user-community needs. These applications can range from enterprise off-the-shelf applications to custom and home-grown applications.

The challenge for security organizations is keeping up with these application developers. In many cases, to expedite delivery of these applications, developers have been known to implement applications on any port that is convenient, or bypass security controls altogether. When they create security backdoors to manage these applications from home or on the road, these can become avenues for attackers to infiltrate.

The ease of application creation and delivery due to virtualization technologies exacerbates the problem. This creates a paradox for security teams that are forced to either be a barrier to business growth by enforcing strict security controls for application delivery or to become helpless to manage security risks in the face of application proliferation.

Rather than attempt to control application developers, the answer lies in controlling the applications. In order to adopt secure application enablement firewall policies, having full and comprehensive visibility into all applications on your network is essential. Understanding how the application landscape has changed is also critical to determine which applications carry threats and whether they should be authorized.

## *Applications Aren't All Good or All Bad*

Applications in the data center can largely be divided into

- ✓ Corporate-supported applications — enterprise off the shelf, custom, and home grown
- ✓ Management applications using RDP, Telnet, and SSH to control the enterprise applications
- ✓ Rogue or misconfigured applications such as peer-to-peer applications for personal use within the data center

The first set of applications described in the preceding list should be allowed for authorized employees, the second set should be enabled only for a select group of IT users, and the third set should be remediated or dropped.

This seems simple enough. However, over the past decade, the application landscape has changed dramatically for organizations. Corporate productivity applications have been joined by a plethora of personal and consumer-oriented applications. This convergence of corporate infrastructures and personal technologies is being driven by mobility and bring your own device (or BYOD) trends, in which organizations are increasingly allowing their employees to use their personal mobile devices — such as smartphones and tablets — in the workplace, for both personal and work-related use.

BYOD isn't just an endpoint challenge. It becomes a data center issue when these personal devices are used to access corporate applications.

Indeed, many organizations now use a variety of social networking applications to support a wide range of legitimate business functions. These functions include recruiting, research and development, marketing, and customer support — and many are even inclined to allow the use of lifestyle applications, to some extent, as a way to provide an “employee friendly” work environment and to improve morale.

Translated into real-world examples in the data center, secure application enablement policies might include allowing

- ✔ IT administrators to use a fixed set of remote management applications (such as SSH, RDP, and VNC) to remotely access servers but blocking their use for all other users.
- ✔ SharePoint server access by various groups within the organization, but restricting mobile access from outside the organization due to the port ranges that must be opened. This is done by potentially enabling external third-party apps to leverage SharePoint web services to directly access SharePoint servers in the data center.
- ✔ ERP applications such as Oracle and SAP, but blocking mobile applications that use nonstandard ports and HTTPS outside the organization to directly access ERP servers.
- ✔ Streaming media applications by category, but applying QoS policies to limit their impact on business VoIP applications.
- ✔ The marketing team using a social networking application such as Facebook to share product documentation with customers, while allowing read access by other users in the organization but blocking posting access.

Today's network security solution in the data center, therefore, must be able not only to distinguish one type of application from the next, but also to account for other contextual variables surrounding its use and to vary the resulting action that will be taken accordingly.

## *Hiding in Plain Sight*

Recent high-profile attacks have shown that cyber threats often use common applications to bypass controls. Once on your network, these threats move with little resistance while hiding in plain sight. When their target has been discovered, exfiltration occurs across known applications such as FTP or an application encrypted with SSL.

Just as an attack or compromise within your physical data center is a significant incident, the impact of a compromise in your virtualized environment is amplified because your workloads, some of which use varied trust levels, and associated data are centralized, without any security barriers between to keep them segmented. If your virtual environment is compromised, the attacker has access to everything. An additional challenge to securing your data center workloads is the fact that security policies and associated updates can't keep pace with the speed of your workload (VM) changes, resulting in a weakening of your security posture.

Existing data center security solutions exhibit the same weaknesses found when they're deployed at a perimeter gateway on the physical network. They make their initial positive control network access decisions based on port information. Then they make a series of sequential, negative control decisions using bolted-on feature sets. There are several problems with this approach:

- **Ports first limits visibility and control.** Their focus on ports first limits their ability to see all traffic on all ports, which means that evasive or encrypted applications, and any corresponding threats that may or may not use standard ports, can slip through undetected.


For example, many data center applications — such as Microsoft Lync, Active Directory, and SharePoint — use a wide range of contiguous ports to function properly. This means you need to open all those ports first, exposing those same ports to other applications or cyber threats.

- **They lack any concept of unknown traffic.** Unknown traffic epitomizes the 80/20 rule — it's a small amount of traffic on every network, but it's high risk. Unknown traffic can be a custom application, an unidentified commercial application, or a threat. Blocking it all — a common recommendation — may cripple your business. Allowing it all is high risk. You need to be able to systematically manage unknown traffic using native policy management tools, thereby reducing your security risks.

## Hidden Lynx: Highlighting the need for east–west protection

In many recent data center breaches, attackers moved laterally across either a physical or virtualized network to accomplish their objectives. In 2013, Symantec documented an attack by the APT group Hidden Lynx. The attackers were able to

gain access to a software supplier’s virtualized environment and move laterally from VM to VM. They first gained access to the network using a SQL injection attack (see the following figure).



Once on the network, they installed Backdoor.Hikit, a Trojan that provides extremely stealthy remote access to compromised systems. The attackers then stole the credentials for the virtual machine that contained the digital code-signing certificates. Using this code-signing infrastructure, the attackers signed 32 malicious files that were found within organizations in the United States Defense

Industrial Base Sector. The signing of these files is significant. The process of digital signing implies trust and therefore simplifies the attackers’ goal of compromising the network.

This case study highlights the level of sophistication attackers have achieved, and clearly demonstrates the need to protect east–west traffic in the data center.

✔ **They require multiple policies, but have no policy reconciliation tools.** Their sequential traffic analysis (stateful inspection, application control, IPS, AV, and so on) requires a corresponding security policy or profile, often-times using multiple management tools. The result is that



your security policies become convoluted as you build and manage a firewall policy with source, destination, user, port, and action, and an application control policy with similar rules, in addition to other threat prevention rules. This reliance on multiple security policies that mix positive (firewall) and negative (application control, IPS, AV) control models without any policy reconciliation tools introduces potential security holes introduced by missed or unidentified traffic.

- ✓ **The security policy update process is cumbersome.** Existing security solutions in the data center do not address the dynamic nature of your cloud environment, and cannot adequately track policies to virtual machine additions, removals, or changes.

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

## *Threats Are Taking a Free Ride*

The increasing prevalence of application-layer attacks is yet another disturbing trend. Threats that directly target applications can pass right through the majority of enterprise defenses, which have historically been built to provide network-layer or port-based protection. Threat developers exploit the same methods (described in the previous section) to infiltrate networks that application developers utilize to promote ease of use and widespread adoption, such as tunneling within applications.

The evasion techniques built into these and many other modern applications are being leveraged to provide threats with “free passage” into enterprise networks. It’s no surprise, therefore, that more than 80 percent of all new malware and intrusion attempts (see Chapter 3) 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.”

## *Management applications*

The weakest security link in data centers is often what's used to manage them. Management application such as SSH, RDP, Telnet, and VNC enable applications to be easily accessed and managed from anywhere, at any time, but can also serve as a threat vector for the data center. IT administrators and application developers require access to the applications, and occasionally these backdoors are either mistakenly left open or proper access control rules are not enforced.

What's interesting is that these types of applications are being used not only by IT, but also by sophisticated employees who want to access their home machine — or someone else's — while they're at work. These intrepid users are accessing their machines and, in the process, exposing themselves and their company to numerous business and security risks.

According to research in the 2015 Palo Alto Networks *Application Usage and Threat Report*, 79 unique remote-access applications were found in use globally across more than 6,600 organizations. More than two-thirds of those organizations had five or more different remote access applications running on their networks. Globally and across all industries, remote access application usage consisted of 48 percent Microsoft Remote Desktop, 16 percent TeamViewer, 11 percent Citrix, 9 percent VNC, 8 percent Telnet, and 8 percent all other applications.

Attackers actively scan for these open backdoors, and when any backdoors are found in a vendor's application, attacks are quickly extended to the vendor's customers and business partners.

The impact from a breach can be catastrophic. Exposed applications could be vulnerable to brute-force password attacks, or the initial access level could enable control of the entire application and become a stepping stone to other intrusions in the enterprise. Remote desktop and management applications must be properly controlled in the data center, and access control must be strictly enforced.

## Ammyy

In recent years, the legitimate remote access application known as Ammyy has commonly been exploited by adversaries in *vishing* (voice phishing) attacks. These attacks have been largely targeted against English-speaking countries and have been fairly successful in duping users into installing the remote access application and giving the adversary access to their systems.

The attack generally starts with a user receiving a phone call from a person purporting to be from Microsoft, Dell, or even their own organization's IT department. The adversary may then claim that the user's system has been discovered to be infected by some form of advanced malware, and the user

must now install a specific application (Ammyy) to remove it. The adversary then directs the user to either the official Ammyy website to download the server software or to another website that hosts the server software. The adversary then asks the user for the code that the Ammyy software generates, giving them complete access to the user's system. At this point, the adversary may claim the malware infection has been fixed or may begin to load actual malware onto the now remotely controlled system to hold the user at ransom or perform other nefarious activities. The industries with the most number of sessions captured for Ammyy usage were the federal government, manufacturing, and energy.

## Unknown applications

Being able to monitor, manage, and control a known application is one thing, but not every application on a network is known and instantly recognized. Most companies accept a variant of the 80/20 rule — most of the traffic is known, but the rest, which is a small amount, is unknown.

Most unknown applications fall into one of three categories:

- ✔ Internal home-grown applications
- ✔ Commercial applications that haven't yet been identified
- ✔ Potential threats

A network security solution in the data center should be able to identify unknown traffic and drill down into specific communications and logs to understand the threat impact. When home-grown applications and commercial applications not previously identified have been characterized and appropriate security policies implemented, it's a reasonable assumption that any other unknown is likely a potential threat.

A data center with zero unknown applications may be a challenging goal, but recognition and active management of the unknowns will go a long way toward reducing the risks of application-enabled threats in the data center.

## Recognizing Data Center Challenges

As discussed earlier, the application landscape has changed. Corporate applications hosted in the data center include a variety of applications, all of which exhibit a variety of characteristics, from port hopping to tunneling.

Organizations need firewall policies that understand business-relevant elements such as application identity, *user identity* (who is using the application), and the types of content or threats embedded within the application.

Using business-relevant elements, you can transform your traditional “allow or deny” firewall policy into a secure application enablement policy. This means more than allowing only what you expressly define and blocking everything else. It means you can build firewall policies that are based on the application itself or an application feature, users and groups, and content, as opposed to port, protocol, and IP address.

This is specifically why traditional network security solutions are not effective in the data center. Port-based rules may allow other applications that should not be allowed in the data center. The strict adherence to relying on port as the initial classification mechanism means that applications directed over nonstandard ports are missed completely, introducing

unnecessary business and security risks. Finding tech-savvy employees using remote access tools on nonstandard ports is not uncommon.

To implement application control, legacy security vendor solutions require that you first build a firewall policy with source, destination, user, port, and action (for example, allow, deny, drop, log). Then, to control applications, you move to a different configuration tab or a separate management application and duplicate information from the firewall policy, adding application and action. Maintaining and reconciling even a small set of firewall and application control policies is challenging. Most medium to large organizations have hundreds — even thousands — of firewall rules, and the multiple policy rule-base approach not only increases the administrative overhead, but it also increases both business and security risks.

## *Data Center Attack Vectors*

The most common Hollywood premise of a data center attack is the physical intrusion in which attackers accomplish the impossible and gain interior access to a data center to disable specific servers or retrieve proprietary information. Today, data centers are located in remote regions, armed with the best physical security systems, with their locations hand-selected based on the propensity to be able to handle natural and man-made disasters. Physical perimeter security incorporates not only chain-link fences, armed guards, and advanced video surveillance systems, but also sophisticated physical access control, including biometrics systems. Actual physical access is also limited to key personnel, so, the reality of an actual physical attack to the data center, while not impossible, is highly unlikely.

In an enterprise data center environment (see Chapter 4), sophisticated modern malware attacks are more common. Modern malware has outpaced traditional antimalware strategies and, in the process, has established a foothold within the enterprise that criminals and nation-states can use to steal information and attack sensitive assets. In particular, initial compromise of a user or asset ultimately leads to a data center

breach because information within the data center is what holds the most promise of financial gain for these attackers.

The rest of this chapter talks about the new threat landscape shaped by modern malware.

# Recognizing Key Characteristics of Advanced Malware

Enterprise information security teams have been doing battle with various types of malware for more than two decades, often ill-equipped with only an arsenal of woefully inadequate signature-based antivirus software. Verizon's 2015 *Data Breach Investigations Report* describes a growing "detection deficit" trend in which the time to compromise and time to detect a breach has diverged over the past decade. Trustwave's 2015 *Global Security Report* found that it takes an average of 188 days from infection to detection of malware "in the wild." That's an awfully long time for an attack — which often begins with a vulnerability exploit or advanced malware infection — to go undetected and, therefore, unmitigated.



A *vulnerability* is a bug or flaw that exists in software and creates a security risk that may be exploited by an attacker. The attacker crafts an *exploit* that targets the vulnerable software, essentially fooling the vulnerable software into performing functions or running code of the attacker's choice.

This poor "catch rate" is due to several factors. Some malware has the ability to mutate or can be updated to avoid detection by traditional antimalware signatures. Additionally, advanced malware is increasingly specialized to the point where an attacker will develop a customized piece of malware that is targeted against a specific individual or network.

Advanced malware leverages networks to gain power and resilience, and can be updated — just like any other software application — so that an attacker can change course and dig deeper into the network, based on what he finds, or to adapt to changes and countermeasures.

This is a fundamental shift compared to earlier types of malware, which were more or less a swarm of independent agents

that simply infected and replicated themselves. Increasingly, advanced malware has become a centrally coordinated, networked application in a very real sense. In much the same way that the Internet changed what was possible in personal computing, ubiquitous network access is changing what is possible in the world of malware. Now, all malware of the same type can work together toward a common goal, with each infected endpoint expanding the attack foothold and increasing the potential damage to the organization.

Here are some important characteristics and capabilities of advanced malware:

- ✔ **Distributed, fault-tolerant architecture:** Advanced malware takes full advantage of the resiliency built in to the Internet itself. Advanced malware can have multiple control servers distributed all over the world with multiple fallback options, and can also potentially leverage other infected endpoints as communication channels, providing a near infinite number of communication paths to adapt to changing conditions or update code as needed.
- ✔ **Multifunctionality:** Updates from command-and-control servers can also completely change the functionality of advanced malware. This multifunctional capability enables an attacker to use various endpoints strategically, in order to accomplish specific tasks such as stealing credit card numbers, sending spam containing other malware payloads (such as spyware), or installing ransomware for the purpose of extortion.
- ✔ **Polymorphism:** A *hash signature* is a cryptographic representation of an entire file or program's source code. Changing just a single character or bit of the file or source code completely changes the hash signature. *Polymorphism* is used to avoid detection by antimalware signatures by regularly mutating to avoid simple hash signature matches. So, polymorphism can produce an infinite number of unique signature hashes for even the smallest of malware programs. Some malware applications have entire sections of code that serve no purpose other than to change the signature of the malware.
- ✔ **Obfuscation:** Advanced malware often uses common obfuscation techniques to hide certain binary strings that are characteristically used in malware and,

therefore, easily detected by antimalware signatures, or to hide an entire malware program. Obfuscation can be implemented using a simple substitution cipher (such as an XOR operation) or more sophisticated encryption algorithms (such as AES), or using a *packer* to compress a malware program for delivery, and then decompress it in memory at runtime.

## *Malware Threats to the Data Center*

Given its flexibility and ability to evade defenses, advanced malware presents an enormous threat to the enterprise. Advanced malware is virtually unlimited in terms of functionality — from sending spam to the theft of classified information and trade secrets. The ultimate impact of advanced malware is largely left up to the attacker, from sending spam one day to stealing credit card data the next — and far beyond, as many cyberattacks go undetected for months or even several years. For example, the Home Depot security breach of 2014 went undetected for five months and resulted in the compromise of more than 56 million payment cards.

### *Targeted intrusions*

Advanced malware is a key component of targeted, sophisticated, and ongoing attacks, and it can be customized to compromise specific high-value systems in a target network. In these cases, an infected endpoint inside the network can be used to steal login credentials and initiate lateral movement in order to gain access to protected systems and to establish backdoors in case any part of the intrusion is discovered. These types of threats are almost always undetectable by traditional signature-based antivirus software. They represent one of the most dangerous threats to the enterprise because they are specifically created with custom components designed to bypass known security vulnerabilities and weaknesses within the targeted organization. These attacks are typically well financed because they'll ultimately yield valuable loot, such as research and development, intellectual property, strategic planning, financial data, and customer information.



## Carbanak: The great bank robbery

Carbanak is one of the latest examples of a targeted attack that began in August 2013 and is currently still active. The attackers have sent spear-phishing emails with malicious CPL attachments or Word documents exploiting known

vulnerabilities. Once inside the victim's network, money is extracted. Each raid has lasted two to four months. To date, the attackers have targeted up to 100 financial institutions, causing aggregated losses estimated at \$1 billion.

## *Advanced persistent threats*

Advanced persistent threats (APTs) are a class of threats that often combine advanced malware and botnet components to execute a far more deliberate and potentially devastating attack. As the name implies, 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.

## New Indicators of Compromise (IoC) for APT group Nitro uncovered

In mid-July 2014, Palo Alto Network Unit 42 identified yet another legitimate website that had been compromised by APT actors and was serving malware. In this case, it was a group commonly referred to as “Nitro,” which was coined by Symantec in 2011. As Unit 42 dug deeper, it found additional compromised legitimate websites and malware from the same group back through March 2014. In most instances, the malware is one commonly referred to as “Spindest,” though “PCClient” and “Farfli” variants are also used by the group.

Historically, Nitro is known for targeted spear-phishing campaigns and using Poison Ivy malware, which was not seen in these attacks. Since at least 2013, Nitro appears to have somewhat modified its malware and delivery methods to include Spindest

and legitimate compromised websites, as reported by Cyber Squared’s TCIRT. Unit 42’s findings indicate that they are continuing to evolve with the addition of PCClient and Farfli variants.

These events impacted at least the following industries, across four waves:

- ✓ A US-based IT Solutions provider
- ✓ The European office of a major, U.S.-based commercial vendor of space imagery and geospatial content
- ✓ A European leader in power technologies and automation for utilities and industry
- ✓ A U.S.-based provider of medical and dental imaging systems and IT solutions

## *Know Thy Enemy*

Hackers today have evolved into bona fide cybercriminals, often motivated by significant financial gain and sponsored by criminal organizations, nation-states, or radical political groups. Today’s cybercriminal has far more resources available to facilitate an attack, has greater technical depth and focus, and is well funded and better organized.

Why is it important to understand who cybercriminals are and what motivates them? Because a hacker sitting in his parents’ basement may be able to break into a corporate

network and snoop around, but he doesn't necessarily know what to do with, say, intellectual property or sensitive personnel data. On the other hand, a rogue nation-state or criminal organization knows all about extortion and exactly what to do or who to sell stolen intellectual property to on the gray or black market. According to Verizon's 2015 *Data Breach and Investigations Report* (DBIR), 96 percent of cyberattacks are motivated by financial gain.

Additionally, criminal organizations and nation-states have far greater financial resources than do independent hackers. Many criminal hacking operations have been discovered, complete with all the standard appearance of a legitimate business with offices, receptionists, and cubicles full of dutiful hackers. These are criminal enterprises in the truest sense, and their reach extends far beyond that of an individual hacker. Cybercriminals today are focused on stealing valuable information. Consequently, it isn't in a cybercriminal's best interests to devise threats that are "noisy" or that are relatively benign. To be successful, a hacker must be fast or stealthy — or both.

For cybercriminals who favor speed over sophistication, their goal is to develop, launch, and quickly spread new threats immediately on the heels of the disclosure of a new vulnerability. The faster a threat can be created, modified, and spread, the better. 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), are unable to keep up — at least during the early phases of a new attack.

This speed-based approach is facilitated by the widespread existence 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 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).

Targeted attacks against specific organizations or individuals are another major concern. In this case, cybercriminals 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. According to Verizon's 2015 DBIR, 85 percent of threat actors targeting organizations are external to the organization.

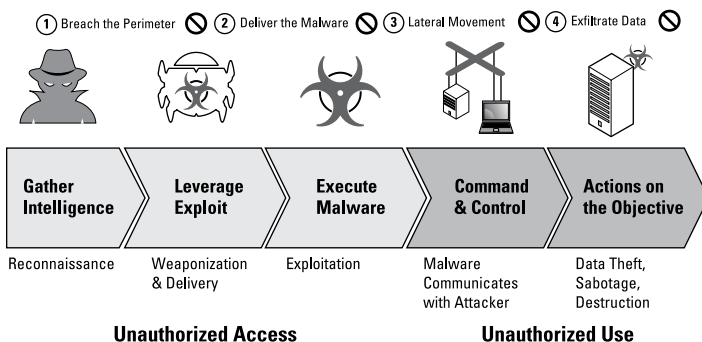
## *Understanding Modern Cyberattack Strategy*

Modern cyberattack strategy has evolved. Instead of a direct, open attack against a high-value server or asset, today's attack strategy employs a patient, multistep covert process that blends exploits, malware, and evasion in a coordinated network attack. The cyberattack life cycle (see Figure 2-1) is a sequence of events that an attacker goes through to successfully infiltrate an organization's network and steal data from it.

Here are the steps of the cyberattack life cycle:

### **1. Reconnaissance.**

Like common criminals, cybercriminals carefully study their victims and plan their attacks. They often use social engineering, phishing, email address harvesting, and other tactics to research, identify, and select targets. They also use various tools to scan networks for vulnerabilities, services, and applications that can be exploited.

**Preventing Access Across the Cyber Attack\* Life Cycle**

\* Intelligence-Driven Computer Network Defense Informed by Analysis of Adversary Campaigns and Intrusion Kill Chains  
Eric M. Hutchins, Michael J. Cloppert, Rohan M. Amin, Ph.D. Lockheed Martin Corporation

**Figure 2-1:** The cyberattack life cycle.

## 2. Weaponization and delivery.

Next, the attacker determines the malware payload and the method that will be used to deliver it. For example, data files or web pages can be weaponized with exploits that are used to target the victim's vulnerable software and delivered via an email attachment or drive-by download.

A *drive-by download* delivers advanced malware or an exploit in the background, usually by taking advantage of a vulnerability in an operating system, web browser, or other third-party application.

## 3. Exploitation.

The attacker generally has two options for exploitation:

- *Social engineering*: A relatively simple technique used to lure someone into clicking a bad link or opening a malicious executable file, for example.
- *Exploits*: A more sophisticated technique in which they essentially trick the operating system, web browser, or other third-party software into running an attacker's code. This means the attacker has to craft an exploit to target specific vulnerable software on the endpoint.

When exploitation has succeeded, an advanced malware payload can be installed.





Using exploits to infiltrate a target network has become an efficient and stealthy method to deliver advanced malware because exploits can be hidden in legitimate files. In addition, readily available off-the-shelf exploit kits significantly reduce the technical knowledge needed to develop exploits. When an exploit is run, the attacker can take control of the endpoint and install malware or run an attack entirely in memory, making it even more difficult to detect because no new files are created on the exploited system.

#### 4. Installation.

When a target endpoint has been infiltrated, the attacker needs to ensure *persistence* (resilience or survivability). Various types of advanced malware are used for this purpose, including the following:

- *Rootkits*: Malware that provides privileged (root-level) access to a computer.
- *Bootkits*: Kernel-mode variants of rootkits, commonly used to attack computers that are protected by full-disk encryption.
- *Backdoors*: Enables an attacker to bypass normal authentication procedures in order to gain access to a compromised system. They're often installed as a failover, in case other malware is detected and removed from the system.
- *Anti-virus software*: May also be installed to disable any legitimately installed antivirus software on the compromised endpoint, thereby preventing automatic detection and removal of malware that is subsequently installed by the attacker. Many anti-virus programs work by infecting the master boot record (MBR) of a target endpoint.

#### 5. Command and control (C&C).

Communication is the lifeblood of a successful attack. Attackers must be able to communicate with infected systems to enable C&C, and to extract stolen data from a target system or network. This communication can also be used by the attacker to target other systems on the victim's network. Thus, the initially infected target may only be the first entry point that

enables lateral movement toward the attacker's ultimate objective. C&C communications must be stealthy and can't raise any suspicion on the network. Such traffic is usually obfuscated or hidden through techniques that include the following:

- *Encryption* with SSL, SSH or some other custom application. Proprietary encryption is also commonly used. For example, BitTorrent is known for its use of proprietary encryption and is a favorite tool — both for infection and C&C.
- *Circumvention* via proxies, remote desktop access tools (such as LogMeIn!, RDP, and GoToMyPC), or by tunneling applications within other (allowed) applications or protocols.
- *Port evasion* using network anonymizers or port hopping to tunnel over open or nonstandard ports.
- *Fast Flux (or Dynamic DNS)* to proxy through multiple infected hosts, reroute traffic, and make it extremely difficult for forensic teams to figure out where traffic is really going.

## 6. Actions on the objective.

Attackers have many different motives for an attack, including data theft, destruction of critical infrastructure, hacktivism, or cyberterrorism. This final phase of the attack often lasts months or even years, particularly when the objective is data theft, because the attacker uses a low-and-slow attack strategy to avoid detection.



The threat landscape has evolved as the use of both business and nonbusiness applications for work-related purposes has become ubiquitous. These “apps” are found in public (for example, SaaS-based applications) and private (for example, virtualized apps in the data center) cloud environments, and are accessed from ever increasing numbers and types of endpoints, including smartphones and tablets. Recognizing that data centers are the “bank vaults” for organizations everywhere, attackers are targeting data centers and the sensitive data in those data centers.





## Chapter 3

# Securing Public and Private Clouds

---

### *In This Chapter*

- ▶ Recognizing security challenges in the cloud
  - ▶ Getting real about network security in virtual environments
- 

**I**n this chapter, you learn about public and private cloud security challenges and how to address them with next-generation firewalls.

## *Public and Private Cloud Security Challenges*

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 in secure segments 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 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 workloads can often be done in minutes, yet 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 need to be identified, and the relevant policy updates determined. In contrast, virtualization 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.

Here are the key requirements for securing the cloud:

- ✔ **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 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 is still deployed in almost every organization, 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. Gartner advocates that organizations "favor security vendors that span physical and virtual environments with a consistent policy management and enforcement framework." In order to ensure that security keeps pace with the speed of change your workflows may exhibit, your security solution should include features that will allow you to lessen, and in some cases, eliminate the manual processes that security policy updates often require.



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
- ✔ Complying and compartmentalizing
- ✔ Keeping pace with the business

Preventing threats has become more difficult in the last several years. Basic attacks on infrastructure have given way to multivector, 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 PCI's Data Security Standards (DSS), U.S. healthcare mandates, and European privacy regulations — are pushing network segmentation deeper into organizations generally, and into data centers and cloud environments specifically.

Finally, maintaining performance and availability usually translates to simplicity. Needless complexity can introduce additional integration issues, outages, and latency. Keeping the data center design and architecture simple is essential.

## *Securing the Cloud with Next-Generation Firewalls*

Today's application and threat landscape renders traditional port-based firewalls and other security solutions largely ineffective at protecting applications and data in public and private cloud environments.

Next-generation firewalls provide key differentiating features, explained in the following sections, to uniquely address these challenges.

## *Ensuring visibility and control*

A next-generation firewall performs a true classification of data and application traffic, based not simply on port and protocol but on 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.

By understanding the types of applications your organization is running in the cloud, you can granularly control the applications that are allowed. And within those applications, you can control the behaviors that are allowed based on a user's role. For example, common applications such as SSH that are used to remotely manage applications need to be tightly controlled, restricted to approved users, and monitored and logged.

This helps in several ways:

- ✔ Enforcing the usage of applications by the appropriate users or groups ensures that business policies are being followed and compliance needs are met.
- ✔ Forcing applications to use standard ports with application policies ensures firewall protection at the application level.
- ✔ Using application visibility and control to prevent or limit high-risk applications reduces the attack surface and limits threats.

This is particularly useful in a cloud environment, where meaningful segmentation by user and application with full content scanning can address compliance requirements and auditability in the cloud. For example, an organization can segment servers containing credit card data and only permit access to that segment for finance users employing a payment application, thus containing and limiting access while maintaining individual accountability. Additionally, threat scanning can be enabled for the segment and data filtering can be enabled to prevent credit card data from flowing out of that segment.

## *Preventing unknown traffic*

Next-generation firewalls also provide a fully integrated approach to threat prevention in a unified context.

This means true coordination of multiple security disciplines — such as application-enabled threat vectors, malware detection, intrusion prevention, file-type controls, and content inspection — for a more intelligent and definitive understanding of threats in the cloud.

The ability to pinpoint and analyze unknown traffic is essential in the cloud, where unknown traffic should ideally constitute a very small percentage of total traffic. A next-generation firewall provides the ability to categorize and analyze unknown traffic to determine whether the traffic is being generated by a legitimate application that isn't recognized or by a potential malware infection.

Unusual traffic patterns such as traffic going to known malware sites, recently registered domains, IP addresses instead of domain names, and the presence of IRC traffic may indicate the presence of advanced malware in a cloud environment. Manual investigation usually means reviewing voluminous log files, much like looking for a needle in a haystack. But next-generation firewalls provide threat logs, automated reports, and easy correlation in a single user interface that eliminates the need to manually track and correlate events. The next-generation firewall threat prevention solution also includes the ability to automatically analyze unknown files in a sandbox environment to identify the behaviors of malware.

## *Protecting users anywhere*

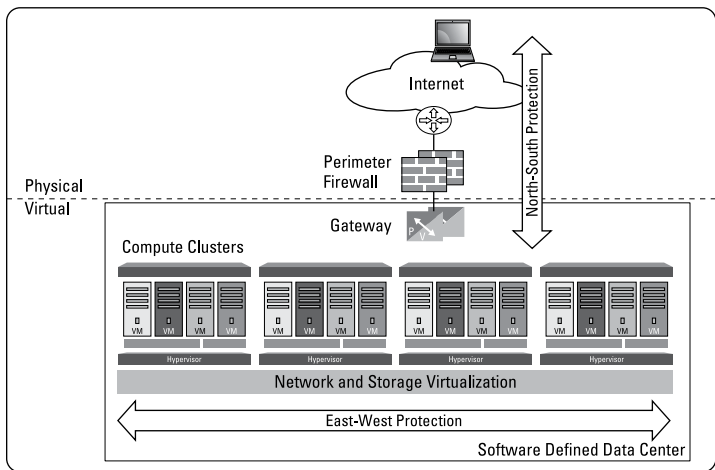
Traditional perimeter-based security is no longer possible with users accessing public and private clouds from practically any device from anywhere in the world. This is a security challenge because traditional security solutions can only enforce security policies when the user is inside the perimeter. A user's behavior also tends to be riskier in the cloud. This behavior increases the likelihood of clicking a dangerous link or visiting a site that serves up a malicious drive-by download.

Next-generation firewalls protect users by enforcing full enterprise firewalling and threat prevention, regardless of device or location. Endpoint protection enables administrators to allow access only after validation of the health or status of the device — such as the latest operating system patches or other settings — required by the organization's security policies.

Finally, access to public and private cloud applications is secured using IPsec or SSL VPN connections.

## Controlling east–west traffic

To better understand the need to secure intra-VM (east–west) traffic, it's helpful to establish an architectural framework. Figure 3-1 displays a typical virtualized data center (private cloud) design.



**Figure 3-1:** Typical virtual data center (private cloud) design architecture.

The compute cluster is the building block for hosting the application infrastructure and provides the necessary resources in terms of compute, storage, networking, and security. Compute clusters can be interconnected using layer 2 or layer 3 technologies such as VLAN, VXLAN, or IP, providing a domain extension for workload capacity. Innovations in the virtualization space allow VMs to move freely in this private cloud while preserving compute, network, storage, and security characteristics and postures.

In a private cloud, there are two different types of traffic, each of which is secured in a different manner:

- **North–south** refers to data packets that move in and out of the private cloud environment. North–south traffic is

secured by one or more perimeter edge firewalls, which control all the traffic into the virtualized data center.

- ✔ **East–west** refers to data packets moving between virtual workloads entirely within the private cloud. East–west traffic is protected by a local, virtualized firewall instantiated on each hypervisor. East–west firewalls are inserted transparently into the application infrastructure and do not necessitate a redesign of the logical topology.

Historically, organizations would implement security to protect traffic flowing north–south, which is insufficient for protecting east–west traffic within a private cloud.



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.

Oftentimes, the question of whether application control is applicable in the cloud arises due to the limited number of known applications that are typically in use. The theory is that we know which applications are in use, so we can more easily secure them. The reality is that recent high-profile breaches have shown that attackers will use applications commonly found on your network and in the cloud to carry out their attacks and extract your data. For example:

- ✔ According to the iSight Partners report on the Target breach, FTP and Webdav were used by attackers to navigate across the network while stealing credit card and user data. This pattern of usage exemplifies how attackers are hiding in plain sight using common applications. Based on Palo Alto Networks 2015 *Application Usage and Threat Report (AUTR)*, these applications are among the top ten applications used to deliver unknown threats.
- ✔ Remote access tools, such as Microsoft RDP, TeamViewer, Citrix, and VNC, are commonly used by attackers to infiltrate networks. According to the 2015 AUTR, a total of 79 different remote access applications were found on more than 7,000 networks worldwide. The AUTR also found that 16.5 percent of files transferred over remote access applications were malicious.
- ✔ Many business applications such as Microsoft Lync, SharePoint, and Active Directory use a wide range of



contiguous ports — including 80, 443, and a range of high number ports — making application control a necessity as a means of allowing only Lync (for example), but no other applications to move across commonly used ports.

- ✔ On average, 8 percent to 10 percent of your network traffic is unknown — it can be an internal application, an unidentified commercial off-the-shelf application, or a threat. The critical functionality you need is the ability to systematically control unknown traffic by quickly analyzing unknowns, determining what it is and where it's coming from, and then managing it through policies, custom applications, or threat prevention profiles.

In each of these examples, next-generation firewalls enable organizations to implement security policies based on Zero Trust principles to help improve their security posture. The concept of Zero Trust extends the practice of network segmentation to the level of granting access based on specific applications, allowing user access based on their credentials and controlling what content can be sent at each segmentation point — all on a never trust, always verify basis. For example:

- ✔ Validate that SharePoint is in use, forcing it over its standard ports and implicitly blocking any other applications from being used.
- ✔ Grant web front-end access to SharePoint over a defined set of ports and applying application specific threat prevention policies.
- ✔ Limit access to the Microsoft SQL database of the SharePoint application itself, implicitly blocking the web front-end from connecting to the database.
- ✔ Allow marketing users, based on their user group membership, to access only SharePoint documents and no other features. Enable only the IT group to use SharePoint Admin, while inspecting the traffic using application-specific threat prevention policies.
- ✔ Identify and block misconfigured or rogue applications like RDP or TeamViewer, leveraging the deny all else premise a firewall follows, or blocking them explicitly with policy.
- ✔ Systematically manage unknown traffic by policy. Create a custom application ID for internal applications, allowing you to control access based on the user and inspect

them for known and unknown malware. Unidentified, commercial applications can be blocked by policy, and submitted for application ID development. Finally, forensics tools and reporting can help you eliminate unknown traffic that may be threat related.

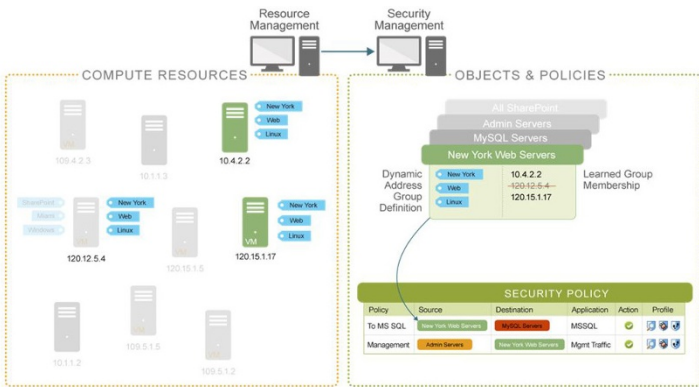
The practice of securing your applications using Zero Trust principles applies to both traditional data centers and cloud computing environments, allowing you to control access based on the application or compute workload and user identity, while blocking potentially rogue or misconfigured applications and preventing any threats from compromising your cloud environment and moving laterally.

## *Keeping pace with the business*

The practice of security has traditionally been methodical, rigorously adhering to well-defined change control and other business processes. When VMs migrate from one physical host to another, they tend to break traditional network security tools that rely on physical and/or network-layer attributes to identify and protect servers and applications. In a virtualized data center, network security solutions need to be capable of ensuring that applicable access control and traffic inspection rules continue to be enforced, regardless of the potential for VMs to move from one physical host to another and from one physical data center to another.

In a data center environment, IP addresses used in policies may need to be updated or created dynamically. Figure 3-2 depicts dynamic policy updates in a virtualized data center where users are creating VMs automatically. Every time a VM is provisioned, or moved, a configuration change on the firewall is needed to ensure the security policy appropriately reflects the updated network topology or VM change.

Next-generation firewalls that bind security policies to VMs automate the process of keeping security policies in sync with the creation of VMs. The dynamic nature of these security objects also enables security policies to be maintained with VM workload movement.



**Figure 3-2:** Security policy automation with dynamic objects.

This can also help address challenges with “VM sprawl” where the number of VMs running in a virtualized environment increases because of the ease of creating a new VM rather than making modifications to an existing one.

Additionally, the network link that is used for VM migrations must be secured. A successful attack against the network can compromise the integrity of VMs as they are migrated or can prevent the VM from being migrated in a timely manner, effectively resulting in a denial of service.

As virtualized data center environments move toward the cloud, the management of a virtualized computing environment must be automated in order for applications and resources to be provisioned on an as-needed basis.

The management and orchestration of a virtualized computing environment is complicated and can involve multiple phases. The first phase includes the provisioning of the virtual server, followed by the provisioning of the networking elements in the virtual environment, such as virtual switches. Finally, the last phase is the security provisioning. In order to properly scale, orchestration software is needed to automate many of these processes. In particular, the speed of the security provisioning can’t slow down the other processes.

Next-generation firewalls offer APIs that enable external orchestration software to connect over an encrypted SSL link for management and configuration. By allowing configuration parameters to be seen, set, and modified as needed, turnkey service templates — where existing security service definitions are defined — can be enabled. This management API enables the proper integration with data center orchestration software so that the security features within the next-generation firewall become part of the data center workflow.

## Chapter 4

# Securing the Hybrid Cloud

.....

### *In This Chapter*

- ▶ Understanding how security needs change as your data center evolves
  - ▶ Recognizing additional security requirements in a hybrid cloud environment
- .....

**A**s organizations transition from a traditional data center architecture to a hybrid cloud, enterprise security strategies must be adapted to support changing requirements in the cloud. This chapter presents a phased plan for implementing security in a hybrid cloud model.

## *A Phased Approach to Security in Virtualized Data Centers*

The following approach to security in the evolving data center — from traditional three-tier architectures to virtualized data centers and to the cloud — aligns with practical realities, such as the inertia behind existing best practices and technology investments, the need for IT security staff to gain experience and build necessary skill sets, and the likelihood that organizations will transform their data centers incrementally rather than in one fell swoop.

This approach consists of four phases.

## ***Phase 1: Consolidating servers within trust levels***

Organizations often consolidate servers within the same trust level into a single virtual computing environment — either one physical host or a cluster of physical hosts. Intra-host communications are generally minimal and inconsequential. As a matter of routine, most traffic is directed “off-box” to users and systems residing at different trust levels.

When intra-host communications do take place, the absence of protective safeguards between these virtualized systems is also consistent with the organization’s security posture for nonvirtualized systems. Live migration features are typically used to enable transfer of VMs only to hosts supporting workloads within the same subnet.

Security solutions should incorporate a robust virtual systems capability in which a single instance of the associated countermeasures can be partitioned into multiple logical instances, each with its own policy, management, and event domains. This enables a single physical device to be used to simultaneously meet the unique requirements of multiple VMs or groups of VMs. Controlling and protecting *inter-host* traffic with physical network security appliances that are properly positioned and configured is the primary security focus.

## ***Phase 2: Consolidating servers across trust levels***

Workloads with different trust levels often coexist on the same physical host or cluster of physical hosts. Intra-host communications are limited, and live migration features are used to enable transfer of VMs only to hosts that are on the same subnet and that are configured identically with regard to routing of VM-to-VM traffic.

Intra-host communication paths are intentionally *not* configured between VMs with different trust levels. Instead, all traffic is forced “off-box” through a default gateway — such as a physical network security appliance — before being allowed to proceed to the destination VM. Typically, this can be accomplished by configuring separate virtual switches with separate physical network interface cards (NICs) for the VMs at each distinct trust level.

As a best practice for virtualization, combining workloads with different trust levels on the same server should be minimized. Additionally, live migrations of VMs should be restricted to servers supporting workloads within the same trust levels and within the same subnet. Over time, and in particular, as workloads move to the cloud, maintaining segmentation based on trust levels becomes more challenging.

### ***Phase 3: Selective network security virtualization***

Intra-host communications and live migrations are “architected” at this phase. All intra-host communication paths are strictly controlled to ensure that traffic between VMs at different trust levels is intermediated either by an on-box, virtual security appliance or by an off-box, physical security appliance. Long-distance live migrations (for example, between data centers) are enabled by combining native live migration features with external solutions that address associated networking and performance challenges.

The intense processing requirements of solutions such as next-generation firewall virtual appliances will ensure that purpose-built physical appliances continue to play a key role in the virtualized data center. However, virtual instances are ideally suited for scenarios where countermeasures need to “migrate” along with the workloads they control and protect.

## *Phase 4: Dynamic computing fabric*

Conventional, static computing environments are transformed into dynamic fabrics (private or hybrid clouds) where underlying resources such as network devices, storage, and servers can be fluidly engaged in whatever combination best meets the needs of the organization at any given point in time. Intra-host communication and live migrations are unrestricted.

This phase requires networking and security solutions that are not only capable of being virtualized but are also virtualization-aware and can dynamically adjust as necessary to address communication and protection requirements, respectively. Classification, inspection, and control mechanisms in virtualization-aware security solutions must not be dependent on physical and fixed network-layer attributes.

In general, higher-layer attributes such as application, user, and content identification are the basis not only for how countermeasures deliver protection but also for how they dynamically adjust to account for whatever combination of workloads and computing resources exist in their sphere of influence.

Associated security management applications also need to be capable of orchestrating the activities of physical and virtual instances of countermeasures first with each other and subsequently with other infrastructure components. This is necessary to ensure not only that adequate protection is provided regardless of the fact that workloads may be frequently migrating across data center locations, but also that it's delivered in an optimal manner.

## *Journey to the Cloud — One Step at a Time*

As organizations implement server virtualization in the data center, the normal progression is toward cloud deployments.



As shown in Figure 4-1, as you move from traditional three-tier architectures to virtualized data centers and the cloud, the same sets of requirements still apply.



**Figure 4-1:** Data center evolution and security requirements.


Hybrid clouds are an aggregation of public clouds and on-premises data centers. The same security challenges and requirements are applicable to hybrid clouds, including

- ✔ Safe application enablement
- ✔ Protection from known and unknown threats
- ✔ Granular segmentation capabilities using Zero Trust principles
- ✔ Automation to address the agility requirements and enable security to keep pace with the business

Additional hybrid cloud security challenges include the following:

- ✔ **Secure connectivity for data in-transit between clouds and on-premises infrastructure, for example using encryption and federated or single sign-on (SSO) authentication:** Figure 4-2 shows a typical hybrid cloud deployment scenario extending the physical data center/private cloud to a public cloud with next-generation firewalls providing IPsec VPN connectivity and visibility, application control, prevention of known and unknown threats, and user-based access control.
- ✔ **Contextual control of data exposure:** Data in the cloud can be either *structured* (application data in

Salesforce.com or spreadsheets stored on Box) or *unstructured* (such as web-based email or documents stored on Google Drive) — both can be a source of improper data shares. To properly protect data in the cloud and ensure regulatory compliance for sensitive data, such as cardholder data (for PCI DSS) and PII, you need security tools that enable you to define granular, context-aware policy controls with the ability to drive enforcement and quarantine users and data before a violation occurs.



**Figure 4-2:** Hybrid cloud deployment scenario with next-generation firewalls.



The “cloud” is not a location but rather a framework of elements — such as automation, orchestration, and service monitoring — that allows disparate processes to be stitched together in a seamless manner. By doing so, the individual elements can be easily replicated and offered on an as-needed basis. Security solutions must provide the appropriate hooks into the orchestration and automation planes in order to not slow down workload creation processes. Because cloud services are extended to multiple organizations with different risk levels, multitenant segmentation is also critical.

## Hybrid Cloud Use Cases


The following sections describe several common hybrid cloud use cases.

## Development and test

Application developers need fungible compute resource to scale on-demand for their application development and, thus, often end up using public cloud resources — without the benefits of enterprise security. This scenario is further complicated by the fact that developers and testers typically need access to enterprise services and data that are located within the enterprise. Moving between the public cloud environment and the on-premises environment presents real challenges. Finally, secure remote access via mobile devices has become an increasingly prevalent requirement.

Figure 4-3 depicts a hybrid cloud deployment for development and test environments that

- ✓ Extends multiple development and test networks from a private data center into the public cloud with the same enterprise security policies
- ✓ Secures enterprise data in transit with IPsec connectivity




**Figure 4-3:** Development and test hybrid cloud use case.

## Capacity augmentation

Many organizations require additional capacity for peak demand scenarios or other seasonal variations. For example, an e-commerce retailer may need burstable capacity for Black Friday and Cyber Monday seasonal sales peaks, or an ERP application may require additional capacity for fiscal year-end processing.

Cloud bursting and application scaling in a public cloud enables organizations to manage peak usage, typically without requiring significant changes to the applications themselves. This scenario keeps the organization's data on-premises in existing enterprise services and databases, while extending the web tier into the public cloud — thereby saving cost and time to build additional capacity.

Figure 4-4 depicts a public cloud being leveraged to burst applications for cyclical spikes in demand by extending the web and application tier to augment on-premises capacity demands.




**Figure 4-4:** Capacity augmentation hybrid cloud use case.

## Limiting “shadow” IT

Enterprise security and compliance mandates are at risk when public clouds are used ad hoc by individual users and lines of business (LOBs). The organization loses visibility and control of its application and data usage in this scenario, as data is moved around without being secured in-transit and potentially stored in a public cloud where it may not be protected or may be inappropriately shared.

Figure 4-5 depicts a hybrid cloud deployment scenario that enables the secure use of public clouds by establishing secure connectivity to workloads for separate LOBs within an organization and applying the same security controls in both private and public clouds.



**Figure 4-5:** Limiting “shadow” IT hybrid cloud use case.



## Chapter 5

---

# Ten (Or So) Evaluation Criteria for Network Security in the Virtualized Data Center

---

**T**his chapter helps you assess network security solutions for your virtualized data center by presenting several important features and criteria for you to consider.

### *Safe Application Enablement of Data Center Applications*

Accurate traffic classification — regardless of ports, protocols, evasive tactics, and SSL encryption — is important in any data center. This is even more critical in a virtualized environment where VMs can be quickly instantiated to support application needs — often without appropriate policies or risk analysis. Another common practice in virtualized 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 cloud security solution needs to 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. Applications must be identified and controlled regardless

of whether the traffic traverses the virtual infrastructure (VM-to-VM traffic) or crosses physical server boundaries. Unknown traffic must also be characterized.

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. This includes more fine-grained and appropriate application functions than simply “allow” or “deny,” such as allow but enforce traffic shaping through 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.

## *Identification Based on Users, Not IP Addresses*

User identification tied to the actual user instead of his IP address is essential in enterprises where users are not only dynamic but also mobile — increasingly accessing data center applications from a variety of mobile devices.

User identification is important to not only get full and accurate visibility of user activity on the network, but (together with application identification) it also can provide appropriate control of applications in the data center. User-based policy control along with log viewing and reporting are key requirements for security in the virtualized data center.

## *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. A fully integrated threat solution that addresses a variety of advanced threats is needed to properly secure the virtualized data center.

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.

## *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 802.1Q and port-based VLANs, but also on the ability to integrate at OSI 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 need the ability to support multiple hypervisor types, such as VMware vSphere, Microsoft Hyper-V, and others, 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. 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 resources from beyond the traditional perimeter of the enterprise.

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

## ***One Comprehensive Policy, One Management Platform***

Most data centers today are hybrid environments, composed of both physical and virtual infrastructures. Network security solutions in the data center will also likely include both physical and virtual solutions.

The network security policy management platform for both physical and virtual data centers must be the same; 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-Readiness*

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 environments, leading to security risks and considerable integration challenges. Automated provisioning of network security capabilities, in line with other orchestration elements of the virtualized data center environment, is essential.

## *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.

However, when applications of different trust levels exist within the same virtual cluster, 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 virtualized 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.



# Glossary

---

**802.1q:** The IEEE standard for VLAN tagging on Ethernet networks.

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

**adware:** Pop-up advertising programs that are commonly installed with freeware or shareware.

**APT:** *See* advanced persistent threat.

**backdoor:** Malware that enables an attacker to bypass normal authentication to gain access to a compromised system.

**bare metal (Type 1) hypervisor:** *See* native hypervisor.

**bootkit:** A kernel-mode variant of a rootkit, commonly used to attack computers that are protected by full-disk encryption.

**bot:** A target machine that is infected by malware and is part of a botnet (also known as a *zombie*).

**bot-herder:** The owner or individual that controls a botnet.

**botnet:** A broad network of bots working together.

**bring your own device (BYOD):** A current trend in which organizations are increasingly allowing their users to bring personal mobile devices — such as smartphones and tablets — into the workplace and connect these devices to the network for both personal and work-related purposes.

**BYOD:** *See* bring your own device.

**cold migration:** A migration process requiring the halting of the original physical server or VM, transfer of associated data, and booting of the VM on the physical host. *See also live migration and warm migration.*

**consumerization:** A current trend in which users increasingly find personal technology and applications that are more powerful or capable, more convenient, less expensive, quicker to install, and easier to use than corporate IT solutions.

**DDNS:** *See* dynamic DNS.

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

**distributed denial-of-service (DDoS):** A large-scale attack that typically uses bots to crash a targeted network or server.

**drive-by-download:** Software, often malware, downloaded onto a computer from the Internet without a user's knowledge.

**dynamic DNS (DDNS):** A technique used to update domain name system (DNS) records in real-time.

**File Transfer Protocol (FTP):** A standard network protocol used to transfer computer files from one host to another over TCP ports 20 and 21.

**FTP:** *See* File Transfer Protocol.

**hosted (Type 2) hypervisor:** A hypervisor that runs within an operating system environment (OSE).

**HTTP:** *See* Hypertext Transfer Protocol.

**HTTPS:** *See* Hypertext Transfer Protocol over SSL/TLS.

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

**Hypertext Transfer Protocol over SSL/TLS (HTTPS):** A secure communication protocol widely used on the Internet.

**hypervisor:** Also known as a virtual machine manager (VMM). Allows multiple "guest" operating systems to run concurrently on a single physical host computer.

**inline mode:** An IDS/IPS mode in which the IDS/IPS is positioned directly in the packet flow and the IDS/IPS can

perform actions (such as block, drop, log, or alert) directly on network traffic. See also *promiscuous mode*.

**Internet Protocol Security (IPsec):** A protocol suite for protecting communications over IP networks using authentication and encryption.

**intrusion prevention system (IPS):** A security appliance or software that detects and prevents known vulnerability exploits.

**IPS:** See intrusion prevention system.

**IPsec:** See Internet Protocol Security.

**LDAP:** See Lightweight Directory Access Protocol.

**Lightweight Directory Access Protocol:** An open standards-based protocol for accessing and maintaining distributed directory services.

**live migration:** A migration process in which a new instance of a VM is created before migrating the existing VM. A live migration doesn't require halting of the VM, transfer of associated data, or a reboot of the VM, and all session information is maintained (stateful). See also *cold migration* and *warm migration*.

**logic bomb:** A program that performs a malicious function when a predetermined circumstance occurs.

**malware:** Malicious code that typically damages or disables, takes control of, or steals information from a computer. Malware includes viruses, worms, Trojan horses, logic bombs, rootkits, bootkits, backdoors, spyware, and adware.

**Microsoft Remote Procedure Call (MS-RPC):** A communications protocol used on MS Windows networks.

**MS-RPC:** See Microsoft Remote Procedure Call.

**native (Type 1) hypervisor:** A hypervisor that runs directly on the host computer's hardware. Also known as a *bare metal hypervisor*.

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

**OSI:** *See* Open Systems Interconnection.

**promiscuous mode:** An IDS/IPS mode in which the IDS/IPS captures a copy of network traffic but cannot block or drop any packets. In this mode, an IDS/IPS can only detect and log or alert. *See also inline mode.*

**QoS:** *See* Quality of Service.

**Quality of Service (QoS):** A measure of the overall performance of a network, typically including availability, bit rate, delay, error rate, jitter, and throughput.

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

**Secure Shell (SSH):** A set of standards and an associated network protocol that allows establishing a secure channel between a local and a remote computer.

**Secure Sockets Layer/Transport Layer Security (SSL/TLS):** A transport layer protocol that provides session-based encryption and authentication for secure communication on the Internet.

**security incident and event management (SIEM):** Security technology that provides real-time analysis of network security alerts.

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

**SIEM:** *See* security incident and event management.

**Simple Mail Transfer Protocol (SMTP):** The Internet standard for email using TCP port 25 (by default).

**SMB:** *See* Server Message Block.

**SMTP:** *See* Simple Mail Transfer Protocol.



**spyware:** A form of malware that's installed on a user's computer, often for the purpose of collecting information about Internet usage or for taking control of a computer.

**SSH:** *See* Secure Shell.

**SSL/TLS:** *See* Secure Sockets Layer/Transport Layer Security.

**stateful inspection:** Also known as dynamic packet filtering; maintains the status of active connections through the firewall to dynamically allow inbound replies to outbound connections.

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

**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.

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

**User Datagram Protocol (UDP):** A connectionless-oriented protocol often used for time-sensitive, low-latency communications that don't require guaranteed delivery.

**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.

**Virtual Private Network (VPN):** An encrypted tunnel that extends a private network over a public network (such as the Internet).

**VLAN:** *See* Virtual Local Area Network.

**VPN:** *See* Virtual Private Network.

**warm migration:** A migration process that does not require halting of the VM, transfer of associated data, or a reboot of the VM, and all session information is maintained (stateful). *See also cold migration and live migration.*

**workload:** The amount of processing to be done by a specified computer. Increasingly synonymous with VM.



THIS COULD BE  
**THE END**  
OF BREACHES

Discover the power of Palo Alto Networks Next-Generation  
Prevention Platform. End-to-end cybersecurity for any business.



See where it all stops: [go.paloaltonetworks.com/TheEnd](https://go.paloaltonetworks.com/TheEnd)

These materials are © 2016 John Wiley & Sons, Inc. Any dissemination, distribution, or unauthorized use is strictly prohibited.

# Restore, control, and protect your applications and data in virtualized data centers!

Virtualization is one of the hottest trends of the past decade and a key enabling technology in virtualized data centers and cloud computing strategies. But the modern application and threat landscape, coupled with sophisticated malware and cyberattacks, has also evolved. In order for organizations to realize the full benefits of virtualization and cloud computing, they must adapt their security architectures to address these new realities in their public, private, and hybrid cloud environments.

- *How virtualization and cloud computing deliver operational benefits — and introduce new security challenges*
- *How modern threats hide in applications — and expose your data to new risks in public, private, and hybrid clouds*
- *Why traditional security controls alone are not enough in virtual data centers — and how to regain visibility and control of all traffic in the data center: north–south and east–west*
- *Where to deploy next-generation firewalls — and how to use their capabilities using a phased approach to protect your virtualized data center and cloud environments*

**Lawrence C. Miller** has worked in information technology in various industries for more than 25 years. He is the co-author of *CISSP For Dummies* and has written more than 60 other *For Dummies* books.




**Open the book and find:**

- Why organizations are adopting public, private, and hybrid cloud environments
- How attackers are hiding in plain sight by exploiting weaknesses in existing data center security architectures
- How to protect corporate applications and data in the public cloud
- What to look for when evaluating network security in the cloud

Go to [Dummies.com](https://www.dummies.com)<sup>®</sup> for more!

FOR  
**DUMMIES**<sup>®</sup>  
A Wiley Brand

 Also available as an e-book

ISBN: 978-1-119-24978-8  
Not for resale

# **WILEY END USER LICENSE AGREEMENT**

Go to [www.wiley.com/go/eula](http://www.wiley.com/go/eula) to access Wiley's ebook EULA.