

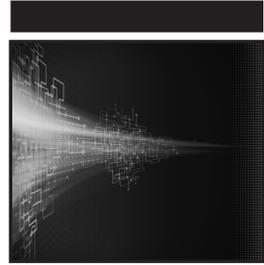


Qing Li
Gregory Clark

SECURITY **INTELLIGENCE**

A Practitioner's Guide to
Solving Enterprise Security Challenges

WILEY



Security Intelligence

A Practitioner's Guide to Solving
Enterprise Security Challenges

Qing Li
Gregory Clark

WILEY

Security Intelligence: A Practitioner's Guide to Solving Enterprise Security Challenges

Published by
John Wiley & Sons, Inc.
10475 Crosspoint Boulevard
Indianapolis, IN 46256
www.wiley.com

Copyright © 2015 by John Wiley & Sons, Inc., Indianapolis, Indiana
Published simultaneously in Canada

ISBN: 978-1-118-89669-3
ISBN: 978-1-118-89667-9 (ebk)
ISBN: 978-1-118-89666-2 (ebk)

Manufactured in the United States of America

10 9 8 7 6 5 4 3 2 1

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 either the prior written permission of the Publisher, or authorization through payment of the appropriate per-copy fee to the Copyright Clearance Center, 222 Rosewood Drive, Danvers, MA 01923, (978) 750-8400, fax (978) 646-8600. 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>.

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 Web site is referred to in this work as a citation and/or a potential source of further information does not mean that the author or the publisher endorses the information the organization or website may provide or recommendations it may make. Further, readers should be aware that Internet websites listed in this work may have changed or disappeared between when this work was written and when it is read.

For general information on our other products and services please contact our Customer Care Department within the United States at (877) 762-2974, outside the United States at (317) 572-3993 or fax (317) 572-4002.

Wiley publishes in a variety of print and electronic formats and by print-on-demand. Some material included with standard print versions of this book may not be included in e-books or in print-on-demand. If this book refers to media such as a CD or DVD that is not included in the version you purchased, you may download this material at <http://booksupport.wiley.com>. For more information about Wiley products, visit www.wiley.com.

Library of Congress Control Number: 2015934208

Trademarks: Wiley and the Wiley logo 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.

*To Huaying, Jane and Adalia;
in Him.*

—Qing Li

*To my parents, James and Mary Clark: thanks for providing guidance early in my
education and career.*

—Greg Clark

*To the cyber security researchers and professionals who are keeping us safe in the
digital world. We offer you our sincere admiration and gratitude for what you do.*

—Qing Li and Greg Clark



Credits

Executive Editor

Carol Long

Project Editor

Rosemarie Graham

Technical Editor

Robert J. Shimonski

Production Editor

Rebecca Anderson

Copy Editor

Marylouise Wiack

**Manager of Content Development
and Assembly**

Mary Beth Wakefield

Marketing Director

David Mayhew

Marketing Manager

Carrie Sherrill

**Professional Technology &
Strategy Director**

Barry Pruett

Business Manager

Amy Knies

Associate Publisher

Jim Minatel

Project Coordinator, Cover

Brent Savage

Proofreader

Kim Wimpsett

Indexer

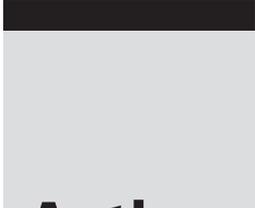
J&J Indexing

Cover Designer

Wiley

Cover Image

©iStock.com/a-r-t-i-s-t



About the Authors

Qing Li is the Chief Scientist and Vice President of Advanced Technologies at Blue Coat Systems, Inc. He is an industry veteran with over 20 years of experience. He has spent the past 11 years designing and developing industry-leading technologies and products at Blue Coat.

Qing is fully responsible for the IPv6 secure proxy, IPv6 WAN optimization technology, and product lines at Blue Coat. He produced the industry's first IPv6 Secure Web Gateway product in 2009 and received the IPv6 Application Solution Pioneer Award from the IPv6 Forum in 2010. Subsequently he produced the industry's first IPv6 WAN Optimization appliance in 2011, and he produced and released the industry's first IPv6 visibility solution in early 2012.

In 2013 Qing took over responsibility for the PacketShaper product. He reinvented the technology and in early 2014 introduced the new PacketShaper S-series appliances into the market place, which are Blue Coat's first 10 Gbps visibility and QoS solutions. The PacketShaper S-series appliance product line reinvigorated new product revenue growth for the first time since 2008, when Blue Coat acquired Packeteer.

In the past five years, Qing's research has concentrated on emerging technologies including advanced application classification algorithms, mobile security, SSL interception, malware detection, and data analytics. His innovations have transformed the Blue Coat technology and product landscape.

Prior to Blue Coat, Qing spent over eight years at Wind River Systems and was the Lead Architect of Wind River's Networking Group. He was responsible for both the pSOS+ and VxWorks networking systems. He led a large distributed team and, in a development partnership with Siemens, successfully delivered VxWorks 6.0 for Network Equipment in early 2003; this was the first VxWorks release that offered full IPv6 support.

Qing is a published author, most notably of a two-volume reference series on IPv6. Volume I, *IPv6 Core Protocols Implementation*, and Volume II, *IPv6 Advanced Protocols Implementation*, were published in 2006 and 2007, respectively, by Morgan Kaufmann Publishers. In 2003 Qing wrote the embedded systems development book *Real-Time Concepts for Embedded Systems*, which was published by CRC Press; it has served as a reference text in the industry as well as in universities. Qing was also a contributing author to *Handbook of Networked and Embedded Control Systems*, a first-of-its-kind book published in 2005 by Birkhäuser.

Qing holds 17 U.S. patents, with many more pending in the areas of networking and security. He has been an active speaker at industry and academic conferences and contributes to discussions of technological innovation and development across a wide range of media around the world.

Gregory Clark is currently the Chief Executive Officer and a member of the Board of Directors of Blue Coat Systems, Inc., a developer of products and services that secure enterprise infrastructure. Mr. Clark previously served as Chief Executive Officer of Mincom, a leading global provider of software and services to asset-intensive industries. Prior to Mincom, he served as Chief Technology Officer and subsequently became President and Chief Executive Officer of E2open, a leader in ERP-agnostic global supply chain integration.

Earlier in his career, Mr. Clark was the IBM Distinguished Engineer responsible for IBM's security technology and served as a vice president at IBM's Tivoli Systems, Inc. Before joining IBM, he founded the security software firm, Dascom, Inc., which was sold to IBM in 1999 and formed a critical element of IBM's security product line. Mr. Clark previously held senior roles with AT&T's UNIX System Laboratories. He is also a member of the Board of Directors of the Global Healthcare Exchange (GHX), Imperva (IMPV), and Emulex (ELX). Mr. Clark is also a Senior Operating Partner at Thoma Bravo. He has almost 30 years of experience in enterprise infrastructure and security and has been granted multiple patents in security technology and business process applications.



Acknowledgments

I want to thank my beautiful wife, Huaying, for replenishing my perseverance with her inexhaustible love and support and for being my best friend and a great mommy. I am blessed with two beautiful girls, Jane and Adalia; they are the joy of my life, my inspiration, and through them I see God's grace. I am also grateful that I can draw my strength from Philippians 4:13, "I am able to do all things in Him who empowers me."

I would like to thank Wenjing Wang and Min Hao Chen for being my research assistants. You guys are simply awesome!

I would also like to thank Chris Larsen, Ron Frederick, Tim van der Horst and Ryan W. Smith for their insightful thoughts. I would like to thank Liliya Bederov for helping with the graphics.

I would like to thank Carol A. Long for recognizing the value of, and being the executive acquisitions editor for, this book. I would also like to thank Rosemarie Graham for her tireless efforts at managing the editing and production phase of the book and for pushing it over the finish line.

—Qing Li



Contents

Foreword		xv
Preface		xvii
Chapter 1	Fundamentals of Secure Proxies	1
	Security Must Protect and Empower Users	2
	The Birth of Shadow IT	2
	Internet of Things and Connected Consumer Appliances	3
	Conventional Security Solutions	5
	Traditional Firewalls: What Are Their Main Deficiencies?	5
	Firewall with DPI: A Better Solution?	9
	IDS/IPS and Firewall	11
	Unified Threat Management and Next-Generation Firewall	14
	Security Proxy—A Necessary Extension of the End Point	15
	Transaction-Based Processing	18
	The Proxy Architecture	19
	SSL Proxy and Interception	22
	Interception Strategies	24
	Certificates and Keys	28
	Certificate Pinning and OCSP Stapling	32
	SSL Interception and Privacy	33
	Summary	35
Chapter 2	Proxy Deployment Strategies and Challenges	37
	Definitions of Proxy Types: Transparent Proxy and Explicit Proxy	38
	Inline Deployment of Transparent Proxy: Physical Inline and Virtual Inline	41
	Physical Inline Deployment	41
	Virtual Inline Deployment	43

	Traffic Redirection Methods: WCCP and PBR	44
	LAN Port and WAN Port	46
	Forward Proxy and Reverse Proxy	47
	Challenges of Transparent Interception	48
	Directionality of Connections	53
	Maintaining Traffic Paths	53
	Avoiding Interception	56
	Asymmetric Traffic Flow Detection and Clustering	58
	Proxy Chaining	62
	Summary	64
Chapter 3	Proxy Policy Engine and Policy Enforcements	67
	Policy System Overview	69
	Conditions and Properties	70
	Policy Transaction	71
	Policy Ticket	73
	Policy Updates and Versioning System	77
	Security Implications	77
	Policy System in the Cloud Security Operation	80
	Policy Evaluation	82
	Policy Checkpoint	82
	Policy Execution Timing	84
	Revisiting the Proxy Interception Steps	86
	Enforcing External Policy Decisions	90
	Summary	91
Chapter 4	Malware and Malware Delivery Networks	93
	Cyber Warfare and Targeted Attacks	94
	Espionage and Sabotage in Cyberspace	94
	Industrial Espionage	96
	Operation Aurora	96
	Watering Hole Attack	98
	Breaching the Trusted Third Party	100
	Casting the Lures	101
	Spear Phishing	102
	Pharming	102
	Cross-Site Scripting	103
	Search Engine Poisoning	106
	Drive-by Downloads and the Invisible iframe	109
	Tangled Malvertising Networks	113
	Malware Delivery Networks	114
	Fast-Flux Networks	117
	Explosion of Domain Names	119
	Abandoned Sites and Domain Names	120
	Antivirus Software and End-Point Solutions – The Losing Battle	121
	Summary	122

Chapter 5	Malnet Detection Techniques	123
	Automated URL Reputation System	124
	Creating URL Training Sets	125
	Extracting URL Feature Sets	126
	Classifier Training	128
	Dynamic Webpage Content Rating	131
	Keyword Extraction for Category Construction	132
	Keyword Categorization	135
	Detecting Malicious Web Infrastructure	138
	Detecting Exploit Servers through Content Analysis	138
	Topology-Based Detection of Dedicated Malicious Hosts	142
	Detecting C2 Servers	144
	Detection Based on Download Similarities	147
	Crawlers	148
	Detecting Malicious Servers with a Honeyclient	150
	High Interaction versus Low Interaction	151
	Capture-HPC: A High-Interaction Honeyclient	152
	Thug: A Low-Interaction Honeyclient	154
	Evading Honeyclients	154
	Summary	158
Chapter 6	Writing Policies	161
	Overview of the ProxySG Policy Language	162
	Scenarios and Policy Implementation	164
	Web Access	164
	Access Logging	167
	User Authentication	170
	Safe Content Retrieval	177
	SSL Proxy	181
	Reverse Proxy Deployment	183
	DNS Proxy	187
	Data Loss Prevention	188
	E-mail Filtering	190
	A Primer on SMTP	191
	E-mail Filtering Techniques	200
	Summary	202
Chapter 7	The Art of Application Classification	203
	A Brief History of Classification Technology	204
	Signature Based Pattern Matching Classification	206
	Extracting Matching Terms – Aho-Corasick Algorithm	208
	Prefix-Tree Signature Representation	211
	Manual Creation of Application Signatures	214
	Automatic Signature Generation	216
	Flow Set Construction	218
	Extraction of Common Terms	220
	Signature Distiller	222

Considerations	225
Machine Learning-Based Classification Technique	226
Feature Selection	228
Supervised Machine Learning Algorithms	232
Naïve Bayes Method	233
Unsupervised Machine Learning Algorithms	236
Expectation-Maximization	237
K-Means Clustering	240
Classifier Performance Evaluation	243
Proxy versus Classifier	247
Summary	250
Chapter 8 Retrospective Analysis	251
Data Acquisition	252
Logs and Retrospective Analysis	253
Log Formats	254
Log Management and Analysis	255
Packet Captures	259
Capture Points	259
Capture Formats	261
Capture a Large Volume of Data	263
Data Indexing and Query	264
B-tree Index	265
B-tree Search	267
B-tree Insertion	268
Range Search and B ⁺ -tree	270
Bitmap Index	272
Bitmap Index Search	273
Bitmap Index Compression	276
Inverted File Index	279
Inverted File	279
Inverted File Index Query	281
Inverted File Compression	282
Performance of a Retrospective Analysis System	283
Index Sizes	283
Index Building Overhead	285
Query Response Delay	286
Scalability	288
Notes on Building a Retrospective Analysis System	289
MapReduce and Hadoop	289
MapReduce for Parallel Processing	292
Hadoop	293
Open Source Data Storage and Management Solution	295
Why a Traditional RDBMS Falls Short	295
NoSQL and Search Engines	296
NoSQL and Hadoop	297
Summary	298

Chapter 9	Mobile Security	299
	Mobile Device Management, or Lack Thereof	300
	Mobile Applications and Their Impact on Security	303
	Security Threats and Hazards in Mobile Computing	304
	Cross-Origin Vulnerability	305
	Near Field Communication	306
	Application Signing Transparency	307
	Library Integrity and SSL Verification Challenges	307
	Ad Fraud	308
	Research Results and Proposed Solutions	308
	Infrastructure-Centric Mobile Security Solution	311
	Towards the Seamless Integration of WiFi and Cellular	
	Networks	312
	Security in the Network	313
	Summary	315
	Bibliography	317
	Index	327



Foreword

It is difficult to unlearn something that was once considered fact; it's against human nature. But unlearning and then reimagining is where we find ourselves in the field of information security today. Think about the changes in how we use technology that have happened over the past decade: the unbounded mobility of workers, the adoption of cloud services, and the rise of nation-state hackers, hacktivists bent on destruction, and cyber-criminal organizations that are run like efficient corporations. These shifts are reshaping our profession daily and challenging yesterday's "best practices."

When I began teaching at Columbia University in the mid-2000s, the term *hacking* conjured up images of disaffected teenagers for most people. How quickly that association has changed. The professionalization of hacking has led to massive loss of intellectual property and the theft of countless personal records. It has destroyed companies, threatened nations, and thrust security into the consciousness of people who would otherwise not be concerned with technology.

So where does a modern security practitioner become grounded in the realities of today's security? This book is a great place to start. Qing Li and Greg Clark have both left a permanent stamp on the security industry and continue to help some of the biggest organizations in the world to protect themselves. This book is a great resource for security professionals and cyber warriors, as Qing and Greg share the knowledge they have accumulated from building products that protect more than eighty percent of the Fortune 500 corporations around the world.

As the chairman of the world's largest security conference, and an academic and practitioner, I can tell you there has never been a more important time for you to read this book. Think of it as a primer for security in modern times,

against modern adversaries. What I have always admired about Qing and Greg is that they are grounded in the practical. This is a book that doesn't speak in absolutes—it respects the dynamic nature of information security. It tackles the hard topics like malware detection, application intelligence, and retrospective analysis. It examines the design of a system that can protect modern *endpoints*, which can be anything from workstations, laptops, phones, and tablets to smart refrigerators, power meters, and yet-to-be-conceived devices in the Internet of Things. It also exposes the power of what is still one of the most important weapons we have in the fight against attackers: the security proxy.

If you are new to information security, this book is a terrific modern primer. If you have been in security for a while, you must approach this book with a simple truth in mind: our industry is having to reinvent itself in the face of modern attacks. Eight-character passwords and a defined network perimeter are a part of our industry's past, not its present or future. Come with an open mind and allow Qing and Greg to reintroduce you to tools you thought you knew in the context of today's sophisticated attacks.

In this new era of security, the authors will take you into the world of malware distribution networks and show you how they play a central role in attacks. You'll also learn how modern techniques like sandboxing, security analytics, and fine-grained application controls can be wielded to protect a modern enterprise.

Information sharing is essential for today's security professional. The content in this book can help invigorate thought on how to build better security solutions. It can also help you come up with more relevant questions to ask in areas where you want to attain clarity.

When security is done right, it is not about lockdown and fear. It is about opening possibilities and liberating business instead of stifling it. In that way, this is a very hopeful book, and I hope you will enjoy reading it as much as I have.

Hugh Thompson, Ph.D.
Los Gatos, CA
December 2014



Preface

The digitization of a prodigious amount of information is intensifying, from health care records and educational backgrounds, to employment history, credit reports, and financial statements. Words like *eBilling*, *eStatements*, and *paperless transactions* have become part of our everyday language. The ever-increasing ability to retrieve this digital information online, combined with both the unremitting compilation of such information to extrapolate personal traits and behavior and the explosion of convenient venues for accessing the Internet, should encourage questions in curious minds: “Just how vulnerable are we to threats against personal privacy?” and “Who is at liberty to scrutinize the vast amounts of private data?”

In recent years, the rapid growth of high-bandwidth network infrastructures accompanied by a dramatic reduction in storage costs serve as the catalysts in the construction and commercialization of various cloud-based services, which are offered to both institutions and individuals. These cloud-based services range from personal online backup storage, content-sharing, and collaboration tools to customer relations management (CRM). These services are easily attainable with affordable prices that will only invigorate adoption and proliferation. Naturally, for security-conscious minds, questions arise as to how penetrable these services are by nefarious entities and, when compromised, how limited in scope the resulting damages will be from a specific breach incurred on the cloud community as a whole.

Utility companies, power plants, air traffic control systems, public transit systems, and others are predominately under digital control. Media coverage of specific cyber-attacks that have targeted these critical infrastructures indicates that the frequency of the attacks is escalating and with rapidly evolving sophistication, and these attacks are incurring more severe damages on their targets. These stories may include enticing details that are suspenseful and

entertaining; however, failure to detect, defend, and remediate these threats will effect monetary catastrophe and endanger the population with unimaginable consequences. So, what mechanisms have been contrived to entrap offenders before they assail us under a camouflage of bit streams?

Branches of government and the armed forces restrict information flow and closely inspect each individual's cyber activities. Similarly, organizations such as health care providers, insurance companies, and financial institutions must comply with certain industry rules and regulations. Many sumptuary laws require exhaustive access logging and retrospective analysis. Mining this voluminous data into a structured representation demands interdisciplinary expertise, through a process that sanitizes the raw data, sieves out the relevant subsets, transforms and normalizes the selection, and applies analytics to seek out patterns. Data mining and analytics are critical components of the security envelope. The flexibility and diversity of queries that can be issued against the extracted knowledge measure the quality of the data mining approach. In the security context, the length of time taken to excavate data determines how quickly active threats can be divulged, imminent attacks revealed, and felicitous resolutions conjured in response, instead of reacting with extemporaneous and ineffective countermeasures.

Security implementation and enforcement begins with us thinking in terms of the end goals. These goals must be expressible in plain language. For example, the thoughts of the CIO of a large enterprise may be as follows:

- When Bob accesses Dropbox, I want to prevent him from uploading any files but permit him to download content from his account between 8 a.m. and 5 p.m., at a rate of no more than 256 Kbps. Bob is not allowed to upload files because he is new to the company and is under a three-month probation period. However, he does have access to sensitive marketing information, and I want to prevent him from sharing such information externally. Bob has permission to download files from Dropbox because his manager utilizes Dropbox for file sharing across a distributed team. Because Dropbox is Bob's main online application, I want to limit Bob's network bandwidth utilization so that Dropbox does not over-consume available network resources.
- When Alice runs the Skype application, I want to log her text chat sessions because she works in a restricted financial environment. Due to SEC regulations and U.S. Treasury mandates, financial institutions must monitor employee transactions and online behavior in order to detect insider sabotage, data theft, or security breaches that originate externally. For these reasons, all of Alice's online activities must be logged and analyzed.
- When users visit websites during work hours, I want to disallow them from accessing sites that are categorized as adult entertainment. I want the content of each website to be analyzed in real-time for adult material, and if any is discovered, I want to terminate that user session immediately and send an alert to HR for coaching the user on company policies.

These security goals seem straightforward, yet a plethora of networking and security technologies is necessary to achieve the desired end results. For example, let us try to translate the first goal into an actual implementation and observe the various networking and security disciplines that are involved.

The prerequisite of implementing the first security goal, at a minimum, includes knowing which user initiated the network traffic, which application is associated with which traffic flows, and which specific application action generated the traffic.

When Bob initiates a Dropbox session to `www.dropbox.com`, the associated traffic that is observed on the network does not contain visible user information such as login name simply because the entire session is encrypted using TLSv1. One way to determine the user information is by examining the source IP address and then querying a directory service such as Active Directory for mapping information between the username and the IP address. This method is unreliable because multiple users could be running on the same host machine that is assigned a single IP address. In other words, if both Bob and Alice are using the same multi-user system for accessing Dropbox, then the IP address-to-username mapping approach will not produce accurate identification. Therefore, the most reliable way of extracting the user information is by examining the actual HTTPS payload.

Because the traffic is encrypted, it is impossible to decipher unless there is a way to plant a device in the communication path; this device would act as the man-in-the-middle (MITM) that can communicate with the user as if it were the server, while at the same time communicating with the server on behalf of the user. Even when the application does not utilize data encryption between its client and server, the art of application classification will be the key to associate data flows to user-initiated application actions, such as file download or file upload commands. The data rate must be measured constantly and must be adjusted according to the desired rate, assuming the data flow has been associated with a specific application command.

So, to summarize, this simple example involves technologies ranging from application classification and authentication protocol to encrypted traffic interception and quality of service management. Yet the example we have just presented is only one aspect of enterprise security, which relates to employee online access behavior and resource usage monitoring, followed by enforcement according to defined policies. Monitoring an employee's online activities involves more than just restricting recreational traffic for productivity gain; more importantly, an employee could be the source of various types of security breaches. For example, an employee could visit a well-known reputable website; however, if the site has been compromised by hackers who have installed malicious URLs to alluring content, the unsuspecting employee may follow a web link and download a malicious piece of code unintentionally, which then turns the employee's computer into a sensor for a malicious botnet.

Security tools that rely on a reputation-based rating system to evaluate the safety level of a website cannot protect users from new dynamic URLs that link to malicious content. The just-described scenario is occurring with increasing frequency due to the ever-growing and evolving lures that entice unsuspecting users into the dark corners of the Internet. The employee's personal information could be stolen. However, if, for example, the employee is a health care worker who may have access to millions of private records, then this private data could be compromised on a massive scale, inflicting unimaginable damages on families and individuals. Unfortunately, public disclosures of such incidents have been made at an alarming rate in recent years.

If a security breach has been detected, postmortem analysis of the various security compromises that encompass the breach is critical in constructing adequate and flexible defense mechanisms against similar attacks in the future. Depending on the severity and level of sophistication of the attack, the analysis process is typically comprised of inspecting terabytes, if not petabytes, of data that may include user transaction logs and raw packet captures. The essence of this *retrospective analysis* is data mining, and the goals are, at a minimum, to identify the victim or victims of the attack, the area of the initial penetration, and the speed of dispersion and propagation, and to analyze the threat DNA against the known attacks. The combination of real-time traffic analysis, correlation of events and response, and data recording and analytics, together with vulnerability management, are loosely termed Security Information (or Incident) and Event Management (SIEM). The maturity and sophistication of a security solution, therefore, can be demonstrated in its effectiveness at translating security requirements, articulated from natural language into actionable and enforceable security policies within that solution.

Our book is designed and written for CISOs, network administrators, solutions architects, sales engineers, security engineers who implement security solutions, and developers who are building new generations of security products. Similar to unraveling a math word problem, this book guides the reader through a deciphering process that translates each security goal into a set of security variables, substitutes each variable into a specific security technology domain, formulates the equation that is the deployment strategy, and then verifies the solution against the original problem by analyzing security incidents and divulging hidden breaches, ultimately refining the security formula iteratively in a perpetual cycle.

Fear not, you do not need a Ph.D. to read this book. We do assume that you have a basic understanding of the TCP/IP protocols, the HTTP protocol, and a high-level conceptualization of SSL/TLS technology.

The book is organized into nine chapters.

Chapter 1, "Fundamentals of Secure Proxies," dissects traditional defense technologies, such as firewalls and IDS and IPS systems, to illustrate the deficiencies in legacy security solutions. The *proxy* technology is described in detail

from the developer's perspective. This chapter then demonstrates the power of proxies by diving into the specifics of how SSL interception is achieved.

Chapter 2, "Proxy Deployment Strategies and Challenges," provides definitions of the various types of proxies in terms of their deployment strategy, accompanied by their advantages and disadvantages. A proxy, being a stateful device, is confronted by various and unpredictable network infrastructure designs. This chapter enumerates the top deployment challenges and offers respective solutions in detail.

Chapter 3, "Proxy Policy Engine and Policy Enforcements," leverages the policy language of a real-world security product to illustrate the essential elements of an effective policy system and demonstrates how various components of a policy are implemented in various stages of the traffic processing path.

Chapter 4, "Malware and Malware Delivery Networks," provides an overview of the types of malware that are active in the wild. The ploys, lures, and schemes fashioned by the attacks are illuminated through actual incidents. Advanced persistent threats (APTs) and other sophisticated strategies such as Stuxnet and Flame have been employed as infiltration and cyber weapons to wage warfare among countries. This chapter sheds light on this topic.

Chapter 5, "Malnet Detection Techniques," describes the algorithms that are applied for detecting suspicious URLs and content that lead to malware infection. Techniques employed for trapping and analyzing malware and suspicious code are fully articulated in this chapter, along with a discussion of open-source analysis tools.

Chapter 6, "Writing Policies," offers meticulous detail on policy design for many common security objectives in enterprise environments.

Chapter 7, "The Art of Application Classification," examines the classification techniques for identifying applications accurately over live traffic in real-time. Knowing what traffic is associated with which application is the first step in applying intelligent control. This chapter elucidates the technical complexities behind this challenging class of security problems that are under active research.

Chapter 8, "Retrospective Analysis," discusses the algorithms and techniques for data logging, storage, management, and mining knowledge, all in the context of security intelligence.

Chapter 9, "Mobile Security," focuses on the new and fast-growing mobile computing world, where security is optional. This chapter discusses the various technical challenges that make designing and building mobile security solutions difficult. With millions of applications available for download, mobile application identification is a formidable challenge. This chapter offers a comprehensive overview of the current active research trends in this new discipline.

There are countless books on firewalls, malware and viruses, cryptography, IDS, IPS, data mining, and many related concepts. However, a book is needed that unifies these concepts, analyzes and compares the various solutions, digests the security problems into succinct requirements, and crystallizes the implementation

strategies that correlate to specific technology and solution categories. This book is the missing manual that teaches you how to assemble all those parts into practical solutions that solve real-world enterprise security challenges.

At a minimum, we hope this book can assist you in turning some of those desultory conversations of acronyms into meaningful discussions on enterprise security.

Fundamentals of Secure Proxies

The evolution of the secure proxy is a reflection of the evolution of the web. The proxy began as a gateway that bridged content that was processed and managed by various information systems, and served that content to the open web during the early days of Internet web construction. The term *web proxy server* was given to this general intermediary to reflect its main duty at the time, namely, translating web requests from the Internet to representations that could be understood and fulfilled by different internal systems, and vice versa.

The web has evolved, expanded, and flourished from a content-centric, information-sharing system into an elaborate ecosystem for commerce, an acculturation establishment for Millennials, and a foundation for modern-day cloud computing. The web browser has become the instrument that unlocks all of the wealth the web offers. The fundamental web protocols and technology, such as HTTP, SSL, HTML, XML, Java, and JavaScript, have been amalgamated into a complex conduit, which faces relentless assaults from nefarious forces that try to subvert it for profit. However, private intellectual properties and confidential data hosted in private and protected networks are accessible through a browser over secure connections across the Internet. The web has also been adopted as a system of portals for managing critical infrastructures at municipal, state, and national levels. Consequently, the user and the browser have become attack vectors for breaching corporate as well as national security.

The web proxy has evolved from a content gateway into an essential security gateway that focuses on users, applications, and content. The security proxy

differs from a generic web proxy in that the secure proxy can interpret and intercept more application protocols than just HTTP. Secure proxies, especially when deployed in enterprise environments, serve as both protectors and enablers so that their user community can benefit from the web while minimizing the risk of being victimized by malware delivery networks.

Security Must Protect and Empower Users

The rise of the Internet becoming the foundation of the new era in commerce, culture, communication, education, entertainment, and technology was invasive, with profound impact on our social behaviors. It is now ubiquitous and is an indispensable element of both professional and personal life. At the time of the Internet boom, even long before the advent of mobile computing, the line between work hours and personal time was indistinguishable. With the introduction and rapid adoption of smart phones and tablet computing, there is no longer a distinction between a personal and a work-related computing device. This situation is particularly true for employees who travel a great deal as part of their job functions. For this mobile workforce, a regular laptop computer is typically installed with both personal software and work-related applications. They work wherever and whenever they can while roaming through airports and hotels. The expansion of both the Internet and affordable residential broadband networks has enabled many employees to work from home. Similar to the mobile workforce, the home computer serves as both a personal entertainment and productivity platform and a professional instrument that performs corporate-related job functions. Both computing paradigms raise a dilemma: a well-formed physical perimeter that isolates and guards the enterprise network with traditional IT governance is nonexistent. This lack of separation of personal, private information from corporate intellectual property and data on the same storage device can be a liability for both the employee and the employer.

The Birth of Shadow IT

Business applications are migrating from locally hosted solutions within the enterprise to a cloud-hosted collaborative model. This transition means enterprise users are accessing business-critical applications through their web browser, over the standard web protocols, using a diverse range of computing devices that may not be owned or managed by the enterprise. Consequently, the traditional security practice of the allow-or-deny-all approach is inadequate in managing today's complex web-oriented computing paradigm.

In today's enterprises, users demand the ability to choose from a vast number of applications that they can utilize to maximize their productivity when performing their duties, while at the same time leveraging those same applications

for personal objectives. Because enterprise IT and network access policies tend to be restrictive, many user-chosen applications may not be authorized for use in an enterprise network due to security risks, such as the type of information the application gathers and transmits to entities that are external to the enterprise. The servers that the application communicates with may also be easily compromised by attacks. For example, many organizations prevent users from running Dropbox for file sharing for fear that company-related confidential documents may be leaked as a result of unintentional but careless actions. Another typical restriction is that users are forbidden from running any application that participates in a peer-to-peer (P2P) network. This prohibition is likely the precipitant of the Digital Millennium Copyright Act that was signed into law in the United States in 1998. From an enterprise perspective, any copyright infringing material that is stored and that transits the enterprise network presents serious legal liabilities and ramifications. Application software may be produced by various publishers that range from large commercial vendors to independent software developers. An enterprise may exclude an application from its permissible list based on the publisher and its reputation.

One of the fundamental evolutions that have taken place in the enterprise IT environment is the emergence and growth of *shadow IT*. Employees' desire to circumvent IT restrictions led to the use of shadow IT. In the previous example, if Dropbox were blocked by IT policies, then employees would find alternative mechanisms and tools to share files, thus resulting in shadow IT usage. Consider the following example: sales engineers (SEs) travel constantly, and they need to share files with other SEs, employees, and their customers. E-mail systems implement file size limits such that large files cannot be transferred over e-mail. Because Dropbox has been blocked, these SEs may experiment exhaustively with Box.com, Wuala.com, Google Docs, Google Drive, TeamDrive, SugarSync, OneDrive, CloudMe, or Amazon Cloud Drive until they find a solution that is capable of penetrating the IT security net.

Internet of Things and Connected Consumer Appliances

The *Internet of Things* (IoT) refers to uniquely identifiable embedded devices that are networked, which are reachable and manageable through the Internet infrastructure. These embedded devices have proliferated and matured beyond just smart sensors to more intelligent applications such as smart building and home automation systems. Google's \$3.2 billion acquisition of Nest in January 2014, followed by Samsung's acquisition of SmartThings in August 2014, offers a glimpse into market developments that are shaping the future of the IoT. Much of this IoT can now be accessed and controlled through applications on popular mobile devices such as the Apple iPhone and iPad and Google's Android-based gadgets. For example, a homeowner can use the ADT Pulse app on their iPad to activate or deactivate their ADT home alarm system, check motion sensors,

and watch live video feeds from various video cameras that have been installed in their home. The Tesla Model S iPhone app allows a car owner to track their car's location or start and stop electrical charging of the vehicle.

The IoT has met little resistance as it has gradually become engrained into our daily lives, in what appears to be almost a seamless integration, because convenience and ease-of-use have replaced security at center stage. Securing the IoT is a complex problem. Two main aspects of defense include protecting the IoT device and securing the access channel. The access channel includes the communication between the device and its peer (commonly known as *machine-to-machine communications* [M2M]), and the communication between the device and its operator. Because it is embedded, the IoT device has limited computing power and resources, which limits the device's ability to run sophisticated software such as a virus scanner. Such an embedded device is typically powered by either a custom operating system (OS) or a special variant of a known OS. An embedded OS generally lacks security software that is commonly found in a desktop OS, for example, antivirus software. At the time of this writing, the popular Apple iOS has been on the market for over seven years, yet antivirus software for the iPhone and iPad is limited in both variety and functionality; more importantly, such antivirus software is rarely installed by iOS users. Considering the iPhone is by definition an embedded device, the prospect of antivirus and anti-malware software finding its way into the iPhone as a standard application seems impossible, at least for the next few years.

Running an embedded OS implies that software patches that fix security vulnerabilities may not be released at a regular interval, if such a practice exists at all. Even when such a firmware patch mechanism exists, in most cases the patch process relies on the user to be diligent in exercising security practices, and such a demand on the general population is simply unrealistic. Therefore, these factors indicate that IoT devices can become popular attack targets and can be compromised with relative ease. Once such an IoT device is hacked, user information may be retrieved and the device can in fact cause physical harm to its owner; for example, a hacker shutting off a smoke detector during a house fire can cause physical injury or damage. These IoT devices can also be turned into zombies and become part of a large botnet, which can be commandeered into participating in a planned distributed denial-of-service (DDoS) attack against another target.

Other types of consumer electronic appliances, such as the Sony PlayStation 4 (PS4) and Internet-ready HDTVs, are network-capable and face security threats similar to those faced by IoT devices. An Internet-ready HDTV may not allow its owner to browse and surf the web; however, it permits its owner to log in to Facebook and update their Facebook status through the built-in application. The Facebook account information could be stolen if the Internet-ready HDTV is hacked. The Sony PlayStation owner can purchase games at the PlayStation

Store. The PlayStation Network user account information includes the account holder's birthday and contains a stored credit card number. The user credential to log in to the PlayStation Network to play multi-player online games can be stolen by an attacker who has compromised the PS4, thus putting the account holder's privacy at great risk.

Conventional Security Solutions

The *security posture* of an organization refers to the role security plays in the organization's business planning and its business operation. The security posture encompasses the design and implementation of a well-defined security plan. The security plan is comprised of technical solutions including technology in terms of software, hardware, and services that can be implemented at end points and within the network. The security plan also includes non-technical aspects: employee education on the importance of security as an essential element of business operations; a definition of policies on employee conduct and behavior that conforms to corporate security governance; a definition of policies for achieving regulatory compliance; and a definition of procedures and guidelines on responding to security incidents, both internally and externally.

In essence, the security posture refers to how an organization views security: as a business enabler or as a hindrance and an inconvenience to its operational efficiency. An organization's security posture dictates its practices of security and determines the effectiveness of its security implementation. In today's information age, the availability and timely accessibility of information are important keys to an enterprise's success. Enterprises strive to foster innovation by harnessing the wealth of information capital available on the Internet, while at the same time maintaining an energized and engaged workforce.

Security should afford users the freedom to explore and harvest the riches of the Internet, and alleviate the fear of becoming victims of cyber threats. Existing threats change and new ones emerge as the web evolves; therefore, security postures cannot remain static for long and need regular assessment. It is essential to have an in-depth knowledge of available security solutions, and an understanding of the strengths and the weaknesses of each solution in order to perform assessments such as vulnerability testing, penetration testing, and standards-based auditing. Understanding security technologies is the key to implementing the layered defense that is now mandatory in securing users and enterprise networks.

Traditional Firewalls: What Are Their Main Deficiencies?

The firewall, the most commonly known and referenced security device, was once the motif of security-related conversations and continues to be an

essential element of any network security design. The traditional firewall is still the first line of defense. However, the growing body of threats have long surpassed the capabilities of the traditional firewall. The security landscape is now cluttered with acronyms such as unified threat management (UTM), deep packet inspection (DPI), intrusion detection system (IDS), intrusion prevention system (IPS), secure web gateway (SWG), web application firewall (WAF), next-generation firewall (NGFW), application intelligence and control (AIC), and many more. These acronyms create the perception that perhaps the security threats are largely under control, yet in reality, adroit, menacing malware crafters flourish in the shadows, and security battles rage on with growing ferocity and intensity. The various technologies that are behind the acronyms add confusion and inundate the security implementers with overlapping solutions. These overlapping solutions obscure the deficiencies in the core technologies, and this lack of clarity results in the construction and deployment of inadequate defenses.

The deficiencies of the traditional firewall lie in its inability to examine the packet payload, especially when content is encrypted. The traditional firewall examines layer-2 (L2) to layer-4 (L4) packet header information, such as source and destination IP addresses, L4 protocol type, and L4 source and destination port information, as depicted in Figure 1-1. A firewall rule can be written to compare any header field or bits against any specific values and can define instructions for the firewall to apply one or more actions accordingly. For example, a firewall rule can state, “If an incoming packet is a TCP connection initiation frame (i.e., the TCP header contains the SYN flag bit), then transmit a TCP RESET frame back to the sender.” Basically, this firewall rule blocks all incoming TCP connection requests.

Here is another example of a firewall policy: “If the source IP address is 10.944.108, the protocol is TCP, and the destination port is 6881, then discard the packet.” TCP port 6881 is commonly used by the BitTorrent program for P2P traffic. Enterprise firewalls block this port to prevent employees from downloading questionable content and consuming valuable network bandwidth. This firewall policy can be problematic in actual deployment. First, the popularity of BitTorrent has enabled its adoption by various organizations for legitimate use, for example, by communities that distribute open source software releases. In such cases, blocking TCP traffic on port 6881 would preclude users from permissible use of BitTorrent and, in some cases, would interrupt the only distribution channel for a specific open source project. Therefore, the content of a specific BitTorrent session, instead of simply the destination port, should determine whether such a session is permitted. However, a traditional firewall does not have the ability to perform content analysis. Second, BitTorrent uses port 6881 when the port is available; otherwise, port 6882 and subsequent ports are tried until an unused port is found. As such,