Qing Li
Gregory Clark

SECURITY INTELLIGENCE

A Practitioner's Guide to Solving Enterprise Security Challenges

WILEY

Contents

List of Tables

List of Illustrations

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 malnet 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 extemporary 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.

CHAPTER 1
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.