Sherlock Holmes and The Case of the Advanced Persistent Threat

Ari Juels
RSA Laboratories
Cambridge, MA, USA
ajuels@rsa.com

Ting-Fang Yen
RSA Laboratories
Cambridge, MA, USA
tingfang.yen@rsa.com

Abstract

An Advanced Persistent Threat (APT) is a targeted attack against a high-value asset or a physical system. Drawing from analogies in the Sherlock Holmes stories of Sir Arthur Conan Doyle, we illustrate potential strategies of deception and evasion available in this setting, and caution against overly narrow characterization of APTs.

Keywords

Advanced Persistent Threats, Sherlock Holmes

1. Introduction

An *Advanced Persistent Threat* (APT), in industry terminology, is a sophisticated, targeted attack against a computing system containing a high-value asset or controlling a physical system. APTs often require formidable resources, expertise, and operational orchestration. Nation states are the most aggressive perpetrators.

Over the past few years, the media have disclosed several successful APTs directed against high-profile targets. The Operation Aurora attacks against Google and a dozens of other companies in 2009 [18] reportedly aimed to tamper with critical source code. They prompted Google to withdraw its operations from mainland China, whose government it identified as the originator. Stuxnet [10], active during late 2009, was a sophisticated worm with an arsenal of four zero-day attacks. It targeted industrial control systems for uranium enrichment in Iran, reportedly with success at the Natanz nuclear facility. In 2011, RSA, The Security Division of EMC, announced a breach directed against its widely used SecurID authentication tokens [7, 23]. The company attributed the attack to an (unnamed) nation-state whose ultimate objective was U.S. military contractors. Victims of APTs are diverse and numerous, however. Organizations spanning more than thirty different business categories are known to have been targeted, including non-profits, sports committees, news media, and the energy and electronics industry [3]. Many more successful APTs undoubtedly remain undisclosed; yet others have gone undetected.

Documented APTs often unfold in a common sequence of steps [8, 23, 14]:

1. *Social engineering:* Employees of the targeted organization receive spear-phishing e-mail, often via compromised partner organizations or social networks. Opening the attachment or clicking on an embedded link causes the employee's computer to become infected.
2. *Command-and-control (C2):* A backdoor is installed on the compromised machine that opens it to remote control.
3. *Lateral movement:* Given a foothold in the targeted organization, the attacker uses stolen credentials, elevated privileges, or exploitation of software vulnerabilities to access other internal machines hosting high-value assets.
4. *Data exfiltration:* The attacker exfiltrates the assets to external sites under the attackers' control. Intermediary hosts inside the organization may serve to gather the targeted data, which is often compressed and encrypted for concealment.

It has become common to view this series of steps as a formula for APTs and craft defenses accordingly. Proposed countermeasures aim to detect targeted phishing e-mails [4], exfiltration of compressed or encrypted data [11], or traffic from Poison Ivy [5], a remote access tool used in several APT campaigns [23, 27].

Our key message here is that APT isn't a vector of attack or playbook of tactics. It's a *campaign*. Attackers aren't bound by a formula or limited to cyber attacks or cyber intelligence. Bribery, physical surveillance, confidence games—anything goes.

There's no formula for APTs; tactics and technologies change. But basic strategies of deception and evasion are durable across time. This paper is meant as an exercise in flexible thinking and a caution against narrow APT characterization. Perhaps the stratagems we describe will even appear in future APTs (if they're not already in use today). For a rich, vivid compendium of clever deceptions, we appeal to a literary classic: the Sherlock Holmes stories of Sir Arthur Conan Doyle.

Organization

This paper is organized by Holmes cases. In each section, we briefly summarize one case and use its major element of deception to characterize a strategy of attack in the APT setting. The attacks we describe are:

- Section 2: The *Red-Headed-League Attack*. In "The Adventure of the Red-Headed League," a

crowd of red-headed job applicants mask the hiring of a red-headed victim. A Red-Headed League attack uses a general event to conceal its target.

- Section 3: The *Speckled-Band Attack*. In "The Adventure of the Speckled Band," an assassin commits a murder in a seemingly unbreachable space. A Speckled-Band Attack is one characterized by unexpected methods of entry.

- Section 4: The *Bohemian-Scandal Attack*. In "A Scandal in Bohemia," Holmes frightens the possessor of a valuable photograph into securing it from destruction, thereby revealing its location. A Bohemian-Scandal Attack simulates a threat to flush out a target.

- Section 5: The *Blue-Carbuncle Attack*. In "The Adventure of the Blue Carbuncle," a thief smuggles a rare gem out of a hotel by stuffing it down the throat of a Christmas goose. A Blue-Carbuncle Attack conceals adversarial activity or stolen data within legitimate or benign-looking context.

Section 6 concludes with a discussion of the implications of these strategms to APT defense.

2. The Red-Headed-League Attack

> From north, south, east, and west every man who had a shade of red in his hair had tramped into the city to answer the advertisement. Fleet Street was choked with red-headed folk...

Summary.

In "The Adventure of the Red-Headed League," Holmes's client is a red-headed pawnbroker named Wilson. Urged by his assistant, Wilson had responded to a newspaper want-ad strangely stipulating red hair as a job requirement. Offering a handsome salary for little work, it attracted a large crowd of applicants. Wilson successfully obtained the job, which involved clerical busywork for a society called "The Red-Headed League." The society then mysteriously vanished.

Holmes discovers a deception targeting Wilson himself. The job advertiser was a criminal gang that included Wilson's assistant. They wanted to lure Wilson away from his shop and tunnel under it to rob the vault of a bank next door. The Red-Headed League was a decoy. The crowd of red-headed applicants formed a cover for the existence of a single targeted applicant.

The deception strategy: Encompass a victim in a general event that conceals a targeted attack.

Example: A red-headed botnet.

Given the vast, often unmanageable, range of threats an enterprise Security Operations Center (SOC) confronts, SOC administrators will often dismiss generalized attacks, e.g., botnets, to focus instead on attacks targeting the enterprise. The RSA 2011 Cybercrime Trends Report notes that 88% of Fortune 500 companies display botnet activity associated with their domains [25]. Botnets have traditionally served as launch points for denial-of-service attacks and spam, and targeted external entities, not infected hosts. For this reason, owners of infected hosts often fail to clean them, a phenomenon that helps fuel botnet growth.

This same self-interested calculus causes SOCs to ignore botnets. As a general phenomenon, a botnet seems largely benign to an enterprise. It's consequently an excellent platform for a Red-Headed-League attack.

To launch a Red-Headed-League APT by means of a botnet, an attacker creates, captures [9], or rents [12] a botnet large enough to include machines within the victim's network. From a general-purpose comandeering of host resources, the botnet can then be redirected as a tool for targeted attack. The command-and-control facility of a botnet enables field updates of bot executables. The attacker can reinstrument bots for lateral movement within the enterprise. Having gained a beachhead, the attacker can bypass the standard first APT step of social engineering and malware infection via spearphishing. (Better still, the attacker can mount a social-engineering attack in parallel as a decoy. Forensics may turn up this obvious targeted attack and thus overlook the lower-profile, still potent botnet.)

Other red-headed attacks.

A spearphishing attack usually exploits an existing trust relationship. E.g., a victim is more likely to click on an e-mail attachment ostensibly from a coworker than one out of the blue. But a well-resourced attacker can slowly and anonymously build more durable and seemingly general trust relationships. Creating or contributing to an open-source software community, for instance, could encourage not just downloading of custom-crafted software, but perhaps even its incorporation into products [28]. The friend-finding feature on social networks can be exploited to trick the victim into contacting users who appear to share similar interests and geographic location [15]. Holding an industry workshop just to pass off an infected USB stick to a targeted attendee as a giveaway or proceedings copy is another such possibility. (The victim might even be invited to give an industry talk at a workshop...)

3. The Speckled-Band Attack

> ...it became clear to me that whatever danger threatened an occupant of the room could not come either from the window or the door. My attention was speedily drawn, as I have already remarked to you, to this ventilator...

Summary.

In "The Adventure of the Speckled Band," Holmes is consulted by a woman about her sister, Julia. After sleeping alone in a locked room, Julia emerged uttering the words,"It was the band, the speckled band!" and died. As there were no signs of violence, the authorities attributed the death to fear and nervous shock.

Holmes uncovers a murder, but no human had entered the locked room. The murder was instead committed by means of a trained, venomous snake (a speckled swamp adder) that travelled through a ventilator and down a bellrope to the bed.

The deception strategy: Breach a security perimeter

through an unconventional (and perhaps undreamed-of) means of ingress.

Example: A speckled robot.

Defenders of enterprise IT resources typically expect breaches to originate remotely or with a human infiltrator or insider. Instead, a robot might physically penetrate the facilities of the targeted organization through, e.g., a ventilation system, an open window, or a package in a mailroom. The robot can gather and report back with intelligence, seize high-value digital assets directly from computers, or tamper with IT assets. A "speckled" robot could be especially effective against air-gapped resources or locked server rooms.

Programmable robots with sensing capabilities (touch, light, sound, magnetism, ultrasound, etc.) and wireless communication are already obtainable (e.g., LEGO Mindstorms [1]) for just several hundred dollars.

Other ventilators and bellropes.

The FBI reportedly eavesdrops on suspects by modifying their cell phones [19]. This is not unlike a Speckled-Band attack: infiltration takes place over an unexpected channel. Similarly, malware can be introduced surreptitiously in an APT setting, e.g., from mobile phones to laptops and desktops upon plug-in (as in the case of digital photo frames with pre-installed viruses and trojans [17]), spreading to nearby devices wirelessly via vulnerabilities in Bluetooth [20], or even embedded in device drivers during manufacturing.

Of course, the appearance of an unbreachable space may itself be an illusion created by the attacker. The attacker can remove its "ventilator and bellrope" after the fact by, e.g., patching the vulnerability that it exploited to gain system entry, revoking its own elevated privileges after compromise, or deleting logs giving evidence of its intrusion.

4. The Bohemian-Scandal Attack

> The alarm of fire was admirably done. The smoke and shouting were enough to shake nerves of steel. She responded beautifully.

Summary.

The King of Bohemia requests Holmes' assistance in recovering a photograph of himself and his former mistress, Irene Adler, to avoid a scandal in his upcoming marriage. The King has already made several fruitless attempts by waylaying Ms. Adler, ransacking her house, and diverting her luggage.

Rather than undertaking an extensive search for the photograph, Holmes simulates a fire in Ms. Adler's house. She hastens to save the photo, which lies behind a sliding panel in her sitting-room. The photograph's secret location is thus revealed.

The deception strategy: Create disturbances or simulate threats to the victim to obtain intelligence about a target resource.

Example: A Bohemian APT.

Recommended responses to APT largely fall into the categories of preventing collateral damage and gathering forensic evidence. The 2010 SANS Forensics Incident Response Summit published a report by Bejtlich et al. [6] with a list of suggested actions to take following a breach. This includes quarantining suspected machines, changing compromised user credentials, configuring alternative infrastructure for necessary services, and deploying additional monitoring at critical servers.

However, many of these actions may in fact facilitate the Bohemian-Scandal attack.

The deployment of additional monitoring in certain parts of the network reveals the location of high-value assets. The quarantine or shutdown of suspect machines, changes to compromised user accounts, or the incorporation of custom intrusion detection rules, reveal the extent of the victim's knowledge about the attack. The provision of alternative computing infrastructure reveals critical services required by the organization's operation. It would not be difficult for an attacker, who has been lurking in the targeted organization for an indefinite amount of time, to deliberately leave traces of his presence to launch this attack.

5. The Blue-Carbuncle Attack

> I was leaning against the wall at the time and looking at the geese which were waddling about round my feet, and suddenly an idea came into my head...

Summary.

In "The Adventure of the Blue Carbuncle," a client turns to Holmes when he discovers a blue gemstone in the crop of a Christmas goose he's brought home.

Holmes immediately recognizes the gem as a precious stone stolen from the Countess of Morcar in a hotel a few days previously. Holmes traces the goose to the thief, the head attendant at the hotel. The thief had succeeded in smuggling the gem out of the hotel by planting it in the goose. (Subsequently lost by the thief, the goose ended up with Holmes's client).

The deception strategy: Conceal unauthorized communication within commonplace objects or activities.

Blue carbuncles in APTs.

In documented APTs, high-value assets are typically exfiltrated by obfuscating the data through compression or encryption, and concealing it among common file transfer protocols such as FTP or HTTP [8, 23]. There is no reason, however, that exfiltration cannot take place over other popular services (e.g., DNS, SMTP, or Skype), or via steganographic techniques that embed data in images or PDF files [16].

Truer to the blue-carbuncle spirit, a patient attacker can extract pieces of data slowly over time (cutting the "gem" into smaller pieces), by leveraging unused fields in TCP/IP packet headers [24] or otherwise piggy-backing upon *existing communications* [29]. Much prior literature exists on the construction of such covert channels [30]. The myriad of creative techniques bot-masters have explored for botnet command-and-control are also applicable here, such as uploading stolen information to sites hosted by fast-flux [13], with dynamically generated domain names [21], or proxied by popular services like Google Translate [2].

An attacker also has the option of gathering stolen information on employee laptops and exfiltrating it at

a later time when the machine is brought outside of the targeted organization, e.g., to an employee's home. The attacker can thereby bypass the organization's network policies, however restrictive they may be. Another option is exfiltration via external WiFi hotspots near the facility containing a breached host.

6. Conclusion

This paper is intended as a thought exercise to broaden general conceptualization of APTs. While documented attacks proceed by a number of well-defined steps, we note that an APT is a *campaign*, and attackers are by no means bound by a formula. On the contrary, a variety of deception techniques are possible, as we have illustrated through Sherlock Holmes stories.

Even as awareness of APT grows, detection remains a challenging problem. An important alternative for defenders may be to give attackers a dose of their own medicine, so to speak. Honeypots, planted vulnerabilities, intentionally leaked documents—these are time-honored defensive deceptions; undoubtedly many more deserve consideration. "Security through obscurity" is another well motivated alternative [26].

Another defensive principle borne out by recent experience is the inevitability of total compromise [22]. Whatever defensive steps we take, the best strategy may always be to remain prepared for the worst case.

7. REFERENCES

[1] LEGO Mindstorms. `mindstorms.lego.com`.

[2] Concealing network traffic via Google Translate. `practicalmalwareanalysis.com/2012/02/14/concealing-network-traffic-via-google-translate/`, 2012.

[3] D. Alperovitch. Revealed: Operation Shady RAT. Threat Research, McAfee. Whitepaper, 2011.

[4] R. M. Amin, J. Ryan, and J. van Dorp. Detecting targeted malicious email using persistent threat and recipient oriented features. *IEEE Security & Privacy*, (99), 2011.

[5] B. Binde and R. McRee and T. J. O'Connor. Assessing outbound traffic to uncover advanced persistent threat. SANS Institute. Whitepaper, 2011.

[6] R. Bejtlich. CIRT-level response to advanced persistent threat. SANS Forensics Incident Response Summit, 2010.

[7] A. W. Coviello. Open letter to RSA customers. `www.rsa.com/node.aspx?id=3872`, 2011.

[8] M. K. Daly. Invited talk: The advanced persistent threat. In *USENIX Large Installation System Administration Conf.*, 2009.

[9] B. Stone-Gross et al. Your botnet is my botnet: analysis of a botnet takeover. In *ACM CCS*, pages 635–647, 2009.

[10] N. Falliere, L. O. Murchu, and E. Chien. W32. Stuxnet dossier. Symantec Security Response, 2011.

[11] T. W. Fawcett. *ExFILD: A tool for the detection of data exfiltration using entropy and encryption characteristics of network traffic*. PhD thesis, University of Delaware, 2010.

[12] G. Hoglund. The shadowy world of the advanced persistent threat and botnets. *SC Magazine*, 2010.

[13] Honeynet Project. Know your enemy: Fast-flux service networks. Technical report, The Honeynet Project and Research Alliance, 2008.

[14] E. M. Hutchins, M. J. Cloppert, and R. M. Amin. Intelligence-driven computer network defense informed by analysis of adversary campaigns and intrusion kill chains. In *Intl. Conf. Information Warfare and Security*, 2011.

[15] D. Irani, M. Balduzzi, D. Balzarotti, E. Kirda, and C. Pu. Reverse social engineering attacks in online social networks. In *Conf. Detection of Intrusion and Malware, and Vulnerability Assessment*, 2011.

[16] S. Katzenbeisser and F. Petitolas. *Information Hiding Techniques for Steganography and Digital Watermaking*. Artech House, Inc., 2000.

[17] R. Lemos. Malware hitches a ride on digital devices. `www.securityfocus.com/news/11499`, 2009.

[18] McAfee Labs and McAfee Foundstone Professional Services. Protecting your critical assets: Lessons learned from "Operation Aurora". Whitepaper, 2010.

[19] D. McCullagh and A. Broache. FBI taps cell phone mic as eavesdropping tool. `news.cnet.com/2100-1029-6140191.html`, 2006.

[20] C. Merloni, L. Carettoni, and S. Zanero. Studying Bluetooth malware propagation: The bluebag project. *IEEE Security & Privacy*, 5(2), 2007.

[21] P. Porras, H. Saidi, and V. Yegneswaran. An analysis of Conficker's logic and rendezvous points. Technical report, Computer Science Laboratory, SRI International, 2009.

[22] R. L. Rivest. Illegitimi non carborundum. Invited keynote talk given at CRYPTO 2011.

[23] U. Rivner. Anatomy of an attack. `blogs.rsa.com/rivner/anatomy-of-an-attack/`, 2011.

[24] C. H. Rowland. Covert channels in the TCP/IP protocol suite. *First Monday*, 2(5), 1997.

[25] RSA, The Security Division of EMC. RSA 2011 cybercrime trends report. Whitepaper, 2011.

[26] P. Swire. A model for when disclosure helps security: What is different about computer and network security. *Journal on Telecomm. & High Tech. Law*, 3:163, 2004.

[27] The SecDev Group and Munk Centre for International Studies. Tracking GhostNet: Investigating a cyber espionage network. 2009.

[28] S. R. Vadalasetty. Security concerns in using open source software for enterprise requirements. SANS Institute, 2003.

[29] L. Wei, M. K. Reiter, and K. Mayer-Patel. Summary-invisible networking: Techniques and defenses. *Information Security*, pages 210–225, 2011.

[30] S. Zander, G. Armitage, and P. Branch. A survey of covert channels and countermeasures in computer network protocols. *IEEE Comm. Surveys and Tutorials*, 9(3), 2007.