

A study on Advanced Persistent Threats

Ping Chen, Lieven Desmet, and Christophe Huygens

iMinds-DistriNet, KU Leuven
3001 Leuven, Belgium
`{firstname.lastname}@cs.kuleuven.be`

Abstract A recent class of threats, known as Advanced Persistent Threats (APTs), has drawn increasing attention from researchers, primarily from the industrial security sector. APTs are cyber attacks executed by sophisticated and well-resourced adversaries targeting specific information in high-profile companies and governments, usually in a long term campaign involving different steps. To a significant extent, the academic community has neglected the specificity of these threats and as such an objective approach to the APT issue is lacking. In this paper, we present the results of a comprehensive study on APT, characterizing its distinguishing characteristics and attack model, and analyzing techniques commonly seen in APT attacks. We also enumerate some non-conventional countermeasures that can help to mitigate APTs, hereby highlighting the directions for future research.

Keywords: advanced threat, APT, sophisticated attacks, cyber security

1 Introduction

Cyber attacks have existed since the adoption of the Internet and have evolved a lot in the past decades, from viruses and worms in the early days to malware and botnets nowadays. In recent years, a new class of threat, the “Advanced Persistent Threat” (APT) has emerged. Originally used to describe cyber intrusions against military organizations, the APT has evolved and is no longer limited to the military domain. As highlighted in several large-scale security breaches [12,15,1,29], APTs are now targeting a wide range of industries and governments.

While APT has drawn increasing attention from the industrial security community, a comprehensive and clear understanding of the APT research problem is lacking. This paper presents the result of a detailed study of the APT phenomenon, and contributes a taxonomy of phases, mechanisms, and countermeasures. In this paper, we first identify the characteristics of APT, and compare it to traditional threats in Section 2. In Section 3, we dissect a typical APT attack into six phases, analyzing the techniques that are commonly used in each stage. We also enumerate various countermeasure that can be applied to defend against APT attacks. In Section 3.2, we provide case studies of four APTs, illustrating the adversaries’ tactics and techniques by applying our presented taxonomy and technical analysis.

2 Definition: what is APT?

APTs frequently made global headlines in recent years, and many feel that this term is overloaded, since different people refer to it as different things. Because so many different opinions of what constitutes an APT exist in the commercial market [2,14,23], a clear definition is needed. In this paper, we adopt the definition given by US National Institute of Standards and Technology (NIST), which states that an APT is [17]:

“An adversary that possesses sophisticated levels of expertise and significant resources which allow it to create opportunities to achieve its objectives by using multiple attack vectors (e.g., cyber, physical, and deception). These objectives typically include establishing and extending footholds within the information technology infrastructure of the targeted organizations for purposes of exfiltrating information, undermining or impeding critical aspects of a mission, program, or organization; or positioning itself to carry out these objectives in the future. The advanced persistent threat: (i) pursues its objectives repeatedly over an extended period of time; (ii) adapts to defenders’ efforts to resist it; and (iii) is determined to maintain the level of interaction needed to execute its objectives”.

This definition provides a good base for distinction between traditional threats and APTs. The distinguishing characteristics of APTs are: (1) specific targets and clear objectives; (2) highly organized and well-resourced attackers; (3) a long-term campaign with repeated attempts; (4) stealthy and evasive attack techniques. We elaborate on each of these characteristics below.

Specific targets and clear objectives APT attacks are highly targeted attacks, always having a clear goal. The targets are typically governments or organizations possessing substantial intellectual property value. Based on the number of APT attacks discovered by FireEye in 2013 [11], the top ten industry vertical targets are education, finance, high-tech, government, consulting, energy, chemical, telecom, healthcare, and aerospace. While traditional attacks propagate as broadly as possible to improve the chances of success and maximize the harvest, an APT attack only focuses on its pre-defined targets, limiting its attack range.

As for the attack objectives, APTs typically look for digital assets that bring competitive advantage or strategic benefits, such as national security data, intellectual property, trade secrets, etc., while traditional threats mostly search for personal information like credit card data, or generically valuable information that facilitates financial gain.

Highly organized and well-resourced attackers The actors behind APTs are typically a group of skilled hackers, working in a coordinated way. They may work in a government/military cyber unit [15], or be hired as cyber mercenaries by governments and private companies [9]. They are well-resourced from both financial and technical perspectives. This provides them with the ability to work for a long period, and have access (by development or procurement) to zero-day vulnerabilities and attack tools. When they are state-sponsored, they may even operate with the support of military or state intelligence.

A long-term campaign with repeated attempts An APT attack is typically a long-term campaign, which can stay undetected in the target’s network for several months or years. APT actors persistently attack their targets and they repeatedly adapt their efforts to complete the job when a previous attempt fails. This is different from traditional threats, since traditional attackers often target a wide range of victims, and they will move right on to something less secure if they cannot penetrate the initial target.

Stealthy and evasive techniques APT attacks are stealthy, possessing the ability to stay undetected, concealing themselves within enterprise network traffic, and interacting just enough to achieve the defined objectives. For example, APT actors may use zero-day exploits to avoid signature-based detection, and encryption to obfuscate network traffic. This is different from traditional attacks, where the attackers typically employ “smash and grab” tactics that alert the defenders.

In Table 1, we summarize the differences between traditional threats and APTs for several attack attributes.

	Traditional Attacks	APT Attacks
Attacker	Mostly single person	Highly organized, sophisticated, determined and well-resourced group
Target	Unspecified, mostly individual systems	Specific organizations, governmental institutions, commercial enterprises
Purpose	Financial benefits, demonstrating abilities	Competitive advantages, strategic benefits
Approach	Single-run, “smash and grab”, short period	Repeated attempts, stays low and slow, adapts to resist defenses, long term

Table 1: Comparison of traditional and APT attacks

3 Attack Model: how does APT work?

APT attacks are meticulously planned, and typically have multiple steps involved. While a specific APT attack may have its unique features, the stages of APT attacks are similar and they differ mostly in the techniques used in each stage. To describe the phases of an APT attack, we adopt a six-stage model based on the concept of an “intrusion kill chain” introduced in [7]. Using such a kill chain model helps to understand threat actors’ techniques in each stage, and provides guidance for defense against APT attacks as well.

3.1 Phases of an APT attack

A typical ATP attack will have the following six phases: (1) reconnaissance and weaponization; (2) delivery; (3) initial intrusion; (4) command and control; (5) lateral movement; (6) data exfiltration.

(1) Reconnaissance and Weaponization Reconnaissance is also known as information gathering, which is an important preparation step before launching attacks. In this stage, attackers identify and study the targeted organization, collecting as much as information possible about the technical environment and key personnel in that organization. This information is often gathered via open-source intelligence (OSINT) tools and social engineering techniques.

- **Social Engineering.** Social engineering refers to psychological manipulation of people into accomplishing goals that may or may not be in the target’s best interest. In cyber attacks, it is often used for obtaining sensitive information, or getting the target to take certain action (e.g. executing malware).
- **OSINT.** OSINT is a form of intelligence collection from publicly available sources, and nowadays it typically refers to aggregating information about a subject via either paid or free sources on the internet. Various information can be collected via OSINT, ranging from the personal profile of an employee to the hardware and software configurations in an organization.

Besides simply grabbing information from the web, attackers may also employ data mining techniques and big data analytics to automatically process the gathered data, in order to produce actionable intelligence. Based on the gathered intelligence, APT actors construct an attacking plan and prepare the necessary tools. In order to be successful, attackers typically prepare various tools for different attack vectors, so that they can adapt tactics in case of failure.

(2) Delivery In this stage, attackers deliver their exploits to the targets. There are two types of delivery mechanisms: direct and indirect delivery. For direct delivery, the attackers send exploits to their targets via various social engineering techniques, such as spear phishing.

Indirect delivery is stealthy. In this approach the attackers will compromise a 3rd party that is trusted by the target, and then use the compromised 3rd party to indirectly serve exploits. A trusted 3rd party can be a supplier of software/hardware used in the targeted organization, or a legitimate website that is frequently visited by the targeted persons (watering hole attack).

- **Spear Phishing.** Spear phishing is a targeted form of phishing in which fraudulent emails only target a small group of selected recipients. It typically use information gathered during reconnaissance to make the attack more specific and “personal” to the target, in order to increase the probability of success. The recipient is lured to either download a seemingly harmless attachment that contains a vulnerability exploit, or to click a link to a malicious site serving drive-by-download exploits [27]. In APT attacks, malicious attachments are used more often than malicious links, as people normally share files (e.g., reports, business documents, and resumes) via email in the corporate or government environment.
- **Watering Hole Attack** The concept of a watering hole attack is similar to a predator waiting at a watering hole in a desert, as the predator knows that the victims will have to come to the watering hole. Similarly, rather

than actively sending malicious emails, the attackers can identify 3rd party websites that are frequently visited by the targeted persons, and then try to infect one or more of these websites with malware. Eventually, the delivery accomplishes when the infected webpages are viewed by victims [18]. The use of watering hole attacks have been seen in several APT campaigns [5,6,10].

(3) Initial Intrusion Initial intrusion happens when the attacker get a first unauthorized access to the target’s computer/network. While the attackers may obtain access credentials through social engineering, and simply use them for “legitimate” access, the typical way for intrusion is executing malicious code that exploits a vulnerability in the target’s computer. The attackers first deliver malicious code in the delivery stage, and then in the intrusion stage gain access to target’s computer when the exploit is successfully executed.

In APT attacks, the attackers often focus on vulnerabilities in Adobe PDF, Adobe Flash and Microsoft Office as well as Internet Explorer. While several APT attacks [12,20] have leveraged zero-day exploits for initial intrusion, many APT attacks also employ older exploits that target unpatched applications.

The initial intrusion is a pivotal phase in an APT attack, since the APT actors establish a foothold in the target’s network in this stage. A successful intrusion typically results in the installation of a backdoor malware. From this point, the threat actors connects to the targets’ network. As a result, network traffic is generated, and file evidences are left on the victims’ computers, which gives defenders the chance to detect an APT in an early phase.

(4) Command and Control Upon successfully establishing a backdoor, APT actors use Command and Control (C2) mechanisms to take control of the compromised computers, enabling further exploitation of the network. In order to evade detection, the attackers increasingly make use of *various legitimate services* and *publicly available tools*.

- **Social networking sites.** The attackers register accounts on various social networking sites, and put control information into blog posts or status messages [16].
- **Tor anonymity network.** Servers configured to receive inbound connections only through Tor are called hidden services. Hosting C2 servers in Tor as hidden services makes them harder to identify, blacklist or eliminate.
- **Remote access tools (RATs)** Although often used for legitimate remote administration, RATs are often associated with cyber attacks [3,28]. A RAT contains two components: a “server” residing on a victim’s endpoint, and a “client” that is installed on the attackers machine. In order to make it work, the “server” component needs to be delivered to the target’s machine first, which is often accomplished via spear-phishing emails.

(5) Lateral Movement Once the communication between the compromised systems and C2 servers is established, threat actors move inside the network, in order to expand their control over the targeted organization, which in turn enables them to discover and collect valuable data. Lateral movement usually

involves the following activities: (1) performing internal reconnaissance to map the network and acquire intelligence; (2) compromising additional systems in order to harvest credentials and gain escalated privileges; (3) identifying and collecting valuable digital assets, such as development plans, trade secrets, etc..

This stage typically lasts a long period, because (1) the attackers want to harvest a maximum of information over a long term; (2) the activities are designed to run low and slow in order to avoid detection. As APT actors move deeper into the network, their movements become difficult to detect. APT actors often utilize legitimate OS features and tools that are typically used by IT administrators, and they may also crack or steal credentials to gain legitimate access, which both make their activities undetectable or even untraceable.

(6) Data Exfiltration The primary goal for an APT attack is to steal sensitive data in order to gain strategic benefits, thus data exfiltration is a critical step for the attackers. Typically the data is funneled to an internal staging server where it is compressed and often encrypted for transmission to external locations under the attackers' control. In order to hide the transmission process, APT actors often use secure protocols like SSL/TLS, or leverage the anonymity feature of Tor network [16].

3.2 Case study of APT attacks

In order to better understand the APT attack model, we studied four APT attacks reported in various sources [12,20,29,10], mapping the attackers' action into our six-stage model. The results are shown in Table 2.

3.3 Countermeasures

Due to the complexity and stealthiness of APTs, there is no single solution that offers effective protection. The current best practice is a wide range of security countermeasures resulting a multi-layered defense. However, due to the specific nature of APTs, some of the existing defense systems need to be reengineered to work in the APT context, hereby requiring additional research. For example, while genetic algorithms have been proved useful for malware detection, their applicability in a large dataset is subject of further study. We elaborate on some defense techniques below.

Security Awareness Training Considering the wide use of social engineering techniques (e.g., spear-phishing emails) in APT campaigns, security awareness training plays an important role in defense. Besides the general best security practices, the training should also provide education about APT attacks. According to an APT awareness study [8], more than half of the industries are not awareness of the differences between APTs and traditional threats, and 67% of respondents report the lack of awareness training relative to APTs.

Traditional Defense Mechanisms Traditional defense mechanisms are necessary since they block known attack vectors, and hence increase the difficulty

Name	Operation Aurora [12]	RAS Breach [20]	Operation Ke3chang [29]	Operation SnowMan [10]
Active Time	June 2009 - December 2009	Unknown - March 2011	May 2010 - December 2013	Unknown - February 2014
Recon. and Weaponization	employees' emails, zero-day exploits, backdoor, and C2 tools	employees' emails, zero-day exploits, trojanized docs, backdoor, RAT	officials' emails, trojanized docs, backdoor, and C2 tools	identify weakness in vfw.org, RAT, backdoor
Delivery	spear phishing (malicious links)	spear phishing (malicious xls file)	spear phishing (malicious zip file)	watering hole attack (compromise & infect vfw.org)
Initial Intrusion	drive-by download (CVE-2010-0249)	xls vulnerability (CVE-2011-0609)	victims open the executable file	drive-by download (CVE-2014-0322)
Command and Control	custom C2 protocol, operating on TCP port 443	Poison Ivy RAT	custom C2 protocol, based on HTTP protocol	ZxShell, Gh0st RAT
Lateral Movement	compromise SCM, and obtain source code	Perform privilege escalation, gather SecureID data	compromise internal systems, collect data	unknown
Data Exfiltration	upload data to C2 servers	compress, encrypt data as RAR files, use FTP for transmission	compress, encrypt data as RAR files	unknown, could be US military intelligence

Table 2: Comparison of different APTs

for APT actors. Common countermeasures that must be used are: patch management, anti-virus software, firewalls, host-based intrusion detection systems (HIDS), network-based intrusion detection systems (NIDS), intrusion prevention system (IPS), Security Information and Event Management (SIEM), content filtering software, etc..

Security awareness training and traditional defense mechanisms do not adequately address APTs. Defenders should combine them with the following state-of-the-art countermeasures that are proposed to mitigate APTs.

Advanced Malware Detection Malware is critical for the initial intrusion. Since APT actors often leverage zero-day exploits or custom-developed evasive tools that bypass traditional defenses, the ability to detect advanced malware is important for defense against APTs. Sandboxing execution is a proven technique for analyzing malware's behavior, which allows defenders to identify unknown advanced malware [19]. As advanced malware may leverage various sandbox-evasion techniques [22] to detect the VM environment, it is important to take these sandbox-evasion techniques into consideration when using sandboxing execution. Also, the research challenge in this area is to perform the malware analysis on-line, and in a non-intrusive fashion.

Event Anomaly Detection Since APT actors use various stealthy and evasive techniques, there is no “known bad” pattern that traditional signature-based defense mechanisms could use. Instead of looking for “known bad” item, an effective APT detection approach is to study normal behavior and search for anomalous activities. Anomaly detection includes the detection of suspicious network traffic and suspicious system activities, or “irregular” clusters of activities (potentially obtained through machine learning). Due to the massive amount of data the need to be analyzed in a reasonable time, anomaly detection typically relates to the research problem of big data analytics. There are several researchers proposing the use of big data analytic for APT detection. In [4], Giura & Wang implemented a large-scale distributed computing framework based on MapReduce to process all possible events, which can be used to detect APT attacks. Liu et. al. [13] proved that analyzing a huge volume of HTTP requests with Hadoop and Lucene can help to quickly uncover potential victims based on a known APT victim.

Data Loss Prevention Since the ultimate goal of an APT attacks is the transmission of valuable data from the target’s network to outside, a fully contextually aware data loss prevention (DLP) solution can be deployed as the last line of defense to protect sensitive data against exfiltration. A DLP solution is a system that is designed to detect and prevent potential data breach by monitoring and blocking sensitive data while in-use, in-motion, and at-rest. It requires the defender to identify its sensitive and critical data first, and define policies and rules in a DLP application for protection. An example research solution is [21].

Intelligence-driven Defense Intelligence-driven defense is not a specific defense solution, it is a defense strategy that leverage the knowledge about the adversaries, and adapt defense based on the gathered intelligence [7]. Since APT actors are determined, and typically launch repeated attacks against the target, defenders can create an intelligence feedback loop, which allow them to identify patterns of previous intrusion attempts, understand the adversaries’ techniques, and then implement countermeasures to reduce the risk of subsequent intrusions.

In Table 3, we summarize the attack techniques and tools that commonly seen in each stage of an APT attack. Additionally, we also identify the countermeasures that can be applied in each stage.

4 Related Work

Existing research on APTs are mostly from industrial security community. Traditional security service providers (e.g., McAfee, Symantec) and emerging APT-focused companies (e.g., FireEye, Mandiant) regularly publish technical reports that document cases of APT attacks [18,1,15,11]. In [26], Thonnard et al. conducted an in-depth analysis of 18,580 email attacks that were identified as targeted attacks by Symantec, and through the analysis, they showed that a targeted attack is typically a long-running campaign highly focusing on a limited number of organizations.

Stages	Attack techniques/tools	Countermeasures
Reconnaissance and Weaponization	OSINT, Social engineering Preparing malware	Security awareness training, Patch management, Firewall
Delivery	Spear phishing, Watering hole attack	Content filtering software, NIDS, Anti-virus software
Initial Intrusion	Zero-day exploits, Remote code execution	Patch management, HIDS, Advanced malware detection
Command and Control	Exploiting legitimate services, RAT, Encryption	NIDS, SIEM, Event Anomaly detection
Lateral Movement	Privilege Escalation, Collecting data	Access control, HIDS, NIDS, Event Anomaly detection
Data Exfiltration	Compression, Encryption, Intermediary Staging	Data Loss Prevention

Table 3: Attack techniques and countermeasures in each stage of an APT attack

There are several articles [24,25] that briefly explained APT attacks and discussed the detection techniques. However, they are not as comprehensive as our presented analysis. As for the countermeasures, several academic researchers proposed the use of big data analytics for APT detection [4,13].

5 Conclusion

APTs are sophisticated, specific and evolving threats, yet certain patterns can be identified in their process. In this paper, we focused on the identification of these commonalities. Traditional countermeasures are needed but not sufficient for the protection against APTs. In order to mitigate the risks posed by APTs, defenders have to gain a baseline understanding of the steps and techniques involved in the attacks, and develop new capabilities that address the specifics of APT attacks. By studying public APT cases and the offerings of the security industry, we presented this broad perspective on APT, which should establish common ground within the security community and provide guidance for further defensive research.

Acknowledgements We want to thank the anonymous reviewers for the valuable comments. This research is partially funded by the Research Fund KU Leuven, iMinds, IWT, and by the EU FP7 projects WebSand, NESSoS and STREWS. With the financial support from the Prevention of and Fight against Crime Programme of the European Union (B-CCENTRE).

References

1. Dmitri Alperovitch. Revealed: Operation Shady RAT. 2011.

2. Richard Bejtlich. What Is APT and What Does It Want? <http://taosecurity.blogspot.be/2010/01/what-is-apt-and-what-does-it-want.html>, 2010.
3. James T. Bennett et al. Poison Ivy: Assessing Damage and Extracting Intelligence. 2013.
4. Paul Giura and Wei Wang. Using large scale distributed computing to unveil advanced persistent threats. *SCIENCE*, 1(3), 2013.
5. Will Gragido. Lions at the Watering Hole – The “VOHO” Affair. <http://blogs.rsa.com/lions-at-the-watering-hole-the-voho-affair/>, 2012.
6. Thoufique Haq and Yasir Khalid. Internet Explorer 8 Exploit Found in Watering Hole Campaign Targeting Chinese Dissidents. 2013.
7. E. M. Hutchins et al. Intelligence-Driven Computer Network Defense Informed by Analysis of Adversary Campaigns and Intrusion Kill Chains. In *Proceedings of the 6th International Conference on Information Warfare and Security*, 2013.
8. ISACA. Advanced Persistent Threat Awareness. 2013.
9. Kaspersky. The Icefog APT: A Tale of Cloak and Three Daggers. 2013.
10. Darien Kindlund et al. Operation SnowMan: DeputyDog Actor Compromises US Veterans of Foreign Wars Website. 2014.
11. FireEye Labs. Fireeye advanced threat report 2013. 2014.
12. McAfee Labs. Protecting Your Critical Assets: Lessons Learned from “Operation Aurora”. 2010.
13. Shun-Te Liu et al. A novel search engine to uncover potential victims for apt investigations. In *Proceedings of the 10th IFIP International Conference on Network and Parallel Computing*, 2013.
14. Mandiant. The Advanced Persistent Threat. 2010.
15. Mandiant. APT1: Exposing One of China’s Cyber Espionage Unit. 2013.
16. Information Warfare Monitor and Shadowserver Foundation. Shadows in the Cloud: Investigating Cyber Espionage 2.0. 2010.
17. NIST. Managing Information Security Risk: Organization, Mission, and Information System View. *SP 800-39*, 2011.
18. Gavin O’Gorman and Geoff McDonald. The Elderwood Project. 2012.
19. M. Zubair Rafique et al. Evolutionary algorithms for classification of malware families through different network behaviors. In *Proceedings of the Genetic and Evolutionary Computation Conference*, 2014.
20. Uri Rivner. Anatomy of an Attack. <https://blogs.rsa.com/anatomy-of-an-attack/>, 2011.
21. Matthew Schmid et al. Protecting data from malicious software. In *Proceedings of the 18th Annual Computer Security Applications Conference*. IEEE, 2002.
22. Abhishek Singh and Zheng Bu. Hot Knives Through Butter: Evading File-based Sandboxes. 2014.
23. Symantec. Advanced Persistent Threats: A Symantec Perspective. 2011.
24. Colin Tankard. Advanced Persistent Threats and how to monitor and deter them. *Network security*, 2011(8):16–19, 2011.
25. Gordon Thomson. APTs: a poorly understood challenge. *Network Security*, 2011(11):9–11, 2011.
26. Olivier Thonnard et al. Industrial espionage and targeted attacks: Understanding the characteristics of an escalating threat. In *Proceedings of the 15th Symposium on Research in Attacks, Intrusions, and Defenses*, pages 64–85. Springer, 2012.
27. TrendLabs. Spear-Phishing Email: Most Favored APT Attack Bait. 2012.
28. Nart Villeneuve and James T. Bennett. XtremeRAT: Nuisance or Threat? 2014.
29. Nart Villeneuve et al. Operation Ke3chang: Targeted Attacks Against Ministries of Foreign Affairs. 2013.