The Unforeseen Cost of Offensive Cyber Capabilities

A Research Paper submitted to the Department of Engineering and Society

Presented to the Faculty of the School of Engineering and Applied Science

University of Virginia • Charlottesville, Virginia

In Partial Fulfillment of the Requirements for the Degree

Bachelor of Science, School of Engineering

James McDowell

Spring 2021

On my honor as a University Student, I have neither given nor received unauthorized aid on this assignment as defined by the Honor Guidelines for Thesis-Related Assignments

Advisor

S. Travis Elliott, Department of Engineering and Society

Signature: James McDowell Date: 05/14/2021

James McDowell

Section I: Introduction

230,000 computers infected, 150 nations affected, and $4,000,000,000 in losses, all from just one strain of malware: WannaCry (Kaspersky 2021). WannaCry was not the first worm, or piece of malware that automatically spreads itself. Nor will it be the last. Even its exploit was not unique. New vulnerabilities, or flaws in software allowing unintended behavior, in common software are discovered weekly, if not daily. These vulnerabilities may be largely unimportant, entirely theoretical, or so critical they result in loss of human life. Your email has them, your web browser has them, your operating system has them, and the only thing left up to anyone's control is whether or not we know about them.

Some cybersecurity experts find it absolutely crucial to aggressively search software for these vulnerabilities, as they cannot be repaired unless they are first found. Security researchers spend their lives analyzing software for issues, and upon finding them, engage in "responsible disclosure," privately informing the software maintainers about the vulnerability so it can be fixed. In an ideal world, there would be nothing wrong with this.

Yet, some software maintainers prefer to ignore the vulnerability in hopes that it will go away if ignored. It won't. Common practice now is to give the software maintainers some period of time to fix the vulnerability and should they fail to do so or request more time, release the vulnerability to the public. This threat looming over the maintainers is believed to give them incentive to fix the vulnerability rather than ignore it.

Even if that wasn't an issue and all software maintainers immediately patched all vulnerabilities, there are still problems. In some cases, the researcher will want credit for having made a discovery, and after a patch is created, the researcher will publicize their findings. In other cases, the software update is analyzed by other researchers or bad actors hoping to identify

the vulnerability itself. Despite the update being released for the software, many businesses and individuals have a tendency to delay applying updates or ignore them entirely.

Thus, even in the best case, the discovery of the vulnerability means that those not actively updating software will be left with software with a new publicly known vulnerability. Had the vulnerability not been discovered, it would still be present in the software, but it would not have been public knowledge. This predicament is not unique to the field of vulnerability research. Security researchers also work to develop open source or commercially available cyber attack toolkits, exploits, and malware to be used by professionals to audit networks and their defenses. Yet while these are intended to be used by professionals, bad actors inevitably gain access to these tools, thus lowering the bar for them to run cyber attacks.

This research seeks to provide an understanding of the costs and benefits of developing well-intentioned offensive cyber capabilities, defined as any tools, techniques, and procedures capable of malicious behavior, regardless of the capability's intended use. But to fully understand these costs and benefits, this research first seeks to understand the actors who play a role in the development and usage of offensive cyber capabilities and the relationships between them. To model and understand these actors and their relationships, this research will employ Actor-Network Theory (ANT).

Section II: Actor-Network Theory

Actor-Network Theory is a social theory in the field of science and technology studies centered around the idea that everything exists in a network of relationships between human and non-human actors. ANT is used to understand both the social and technical aspects of these networks, their construction, and their transformation (Hedström, Dhillon & Karlsson 2010), and

one of the core tenets of actor-network theory is that these social and technological aspects are inseparable. ANT also serves as a methodology, in which the researcher traces the actors and their relationships to better understand the network as a whole (Walsham 1997). In this particular case, ANT will be used to model a cyber threat network of relationships between threat actors (those responsible for cyber incidents), developers of offensive cyber capabilities, offensive cyber capabilities themselves, system security solutions, those responsible for information systems' security, and those relying upon information systems' security. In practice, the full network would include many, many more actors, but they are outside the scope of this paper.

Section III: Actors and Their Relationships

Cyberattacks attributed to nation-states are on the rise, with 36% of attacks on North American companies attributed to foreign nations (O'Malley 2020). Given the large percentage of incidents attributed to governments and the extreme levels of sophistication and persistence in their attacks, understanding their role in this network is of paramount importance. Nation-states serve three important roles in the cyber threat network - they develop offensive capabilities, perform cyberattacks, and must fend off cyberattacks. Nation-state-developed capabilities tend to be highly sophisticated, and their capability development includes vulnerability research, exploit creation, or malware development (FireEye 2014). The resulting tools, techniques, and procedures are rarely publicized; instead, they end up being used to support the nation's cyber operations. This is not always the case; both Great Britain's GCHQ (NCSC 2019) and the NSA (NSA 2020) have publicly disclosed vulnerabilities they discovered to give everyone an opportunity to patch the vulnerabilities before they become commonly exploited.

Armed with their custom-developed, highly sophisticated offensive cyber capabilities, and the resources of a whole nation, nation state cyber operators make some of the most dangerous threat actors out there (Knowles 2020). Nation state led operations are the most sophisticated operations around, often lasting years before they are ever discovered. The scale of nation-state threat actors' goals matches the scale of their capabilities. Some target political and diplomatic institutions to influence standings on a geopolitical scale. Others target entire sectors to bring a competitive advantage to institutions based in their nation (Saminottawa 2020), or in some cases, simply to make money (Tsing 2019). Nation-state actors have even used their capabilities to target activists (Schectman & Bing 2019). Finally, nations have used their capabilities as cyber weapons, destroying centrifuges in nuclear facilities in Iran and taking out Ukraine's power grid.

Nation state actors have relationships to capabilities beyond what they themselves have developed. According to Accenture's 2020 threat intelligence report, major advanced persistent threats (APTs) have begun transitioning to use open source capabilities when possible (Accenture 2020). This gives them a number of significant benefits. For actors with less developed cyber capabilities, open source tooling and exploits significantly lower the bar to entry. Rather than having to develop a whole robust toolchain from scratch, actors can use off-the-shelf malware, exploits and command and control (C2) servers at no cost. Note that while some off-the-shelf tooling such as Cobalt Strike requires a license to use, threat actors use "cracked" versions of these tools, bypassing the need for a license. Training operators to use off-the-shelf tooling is also easier given that training resources are available freely on the Internet (Higgens 2017). Additionally, using open source tooling for reconnaissance means that if they get caught, the APT doesn't also end up getting their tooling caught. Finally, if an analyst

catches an APT using open source malware, it is significantly more challenging to identify the actor responsible. Given that many targets are political, if an attack is properly attributed, it can have strong negative diplomatic repercussions. Even if a nation-state doesn't use open source tooling directly, it's not uncommon for analysts to identify customized versions of openly available tooling. For all of these reasons, the relationship between nation state threat actors and open source tooling is a close and well-established one.

As governments are responsible for providing for their citizens' security, many seek to provide defensive guidance and assistance in addition to ensuring their own security. For example, the United States government releases Security Technical Implementation Guides (DISA 2021), NIST standards, and Secure Host Baselines (NSA 2015) that end up defining industry best practices. This relationship is an incredibly important and complex one. As industry best practices get updated, attackers evolve their strategies to combat the new best practices. This leads to defensive organizations adjusting their best practices. And the cycle continues ad infinitum. Publicly available tooling serves to accelerate this cycle. Tools to automatically apply best practices accelerate the application of best practices, and tools to counter these accelerate the need for updated best practices. This cycle will be revisited later in this paper.

Corporations are the next major actor in the network, though they play a much more defensive role than nation-states. That doesn't mean corporations don't develop offensive cyber capabilities; to the contrary, in fact, they develop some of the more sophisticated publicly available tooling and vulnerability research. A number of threat intelligence companies such as CheckPoint will perform their own vulnerability research (CheckPoint 2021). As the vulnerabilities they find get patched, their customers become more secure, and the company's reputation improves. Other cybersecurity guidance companies such as FireEye take it a step

further than vulnerability research and develop their own offensive tooling. This custom tooling allows them to emulate an adversary with novel capabilities and more accurately audit a system's security posture (FireEye 2021). Some large companies also have vulnerability research teams that simply seek to make the Internet a safer place, such as Google's Project Zero (Project Zero 2021). Software maintainers also perform vulnerability research on their own products in an attempt to ensure their customers are as secure as possible (MSVR 2021). Finally, some business models, such as those employed by Rapid7 and Strategic Cyber LLC involve developing threat emulation software such as the Metasploit Framework or CobaltStrike and providing this software to customers. In fact, Rapid7's Metasploit Framework is open sourced, allowing anyone to access, use, and build upon it for free. The relationships between these capabilities and other actors will be explored later in this paper.

Corporations also can work in an offensive role, even if doing so is almost always for a defensive purpose. There are exceptions; some companies such as Ticketmaster have historically ignored the law in hopes of gaining an advantage over competitors (O'Donnell 2021). Further, there is considerable evidence that China-based corporations have hacked Western companies to steal trade secrets, as seen with Huawei's hack of Cisco, among others businesses (DoJ 2020). That said, the majority of corporate offensive cyber operators work as a "red team" or penetration testers. In both roles, the operator is performing offensive actions against a network in an attempt to evaluate a network's defenses and provide recommendations for improvement. A penetration test is an exercise where professional hackers attempt to break into systems in every way possible and identify all vulnerable points of entry. Meanwhile, a red team engagement is an exercise in which professional hackers attempt to emulate a real threat, generally only seeking a few ways to gain initial access and focusing much more on testing the automated defenses of the

network and determining what the network's defenses are capable of detection. The key difference between the two is that a penetration test primarily seeks to identify how a network can get compromised while a red team engagement primarily seeks to find how well a network's defenses can identify and stop a compromise (Kim 2018).

These penetration tests and red team engagement are a major use-case of well-intentioned offensive cyber capabilities. First consider a penetration test. In order to exhaustively enumerate all methods an attacker could use to compromise a network, a penetration testing team needs to be able to act as an adversary might, meaning they would need the ability to try to exploit potential vulnerabilities in the network. Without publicly accessible (or custom built) exploits, a penetration testing team would be unable to effectively audit a network, meaning a dedicated attacker with exploits on hand would likely have a much easier time breaching the network. The same goes for red team engagements; without tooling that can evade detections the way an adversary would, a red team would be unable to determine how well a network's defenses would be able to respond to a real adversary.

Most companies are not built around providing security as a service; their concern regarding security is how secure *they* are, how likely they are to be compromised, and in the case of software companies, how likely their software is to be vulnerable. Generally speaking, companies like this serve two roles in the cyber threat network. They are targets for attackers, and they undergo penetration tests and red team engagements, either from internal teams or outside contractors. Threat intelligence companies like FireEye interact heavily with non-cybersecurity companies to provide security guidance and help ensure they're secure. These security guidance companies often develop their own anti-malware product to automatically provide defense for their clients' network (FireEye 2021). These anti-malware products need to

be able to tell what's good from what's bad - yet another case where well-intentioned offensive capabilities can be used to improve security.

As companies began to prioritize security, many launched so-called "bug bounty" programs, where anyone who discovers a vulnerability in their software will receive a financial reward for discovering and properly disclosing it (HackerOne 2021). This has led to a number of hobbyists pursuing vulnerability research in their free time, and some even making a living off of bug bounties (Zorz 2020). Hobbyists have also developed extremely popular open source malware such as Mimikatz (Greenberg 2017) as well as open source command and control frameworks for managing malware on a system such as SILENTTRINITY and Merlin (Villarreal 2019). These tools are useful beyond just penetration tests or professional use; their availability provides an opportunity for the next generation of cybersecurity experts to learn. They also are used by researchers to better understand what exactly attackers can do, how they can do it, and how it can be automatically detected.

The sheer number of open source offensive cyber capabilities developed by individuals and hobbyists is staggering, and their popularity and prevalence has opened the door to a new kind of threat: so-called "script kiddies." Script kiddies is the name given to individuals with little technical knowledge or offensive skill who try to carry out cyber attacks. In an ideal world, this would never be possible, but with how readily available and easy to use cyber attack platforms there are, the barrier for entry to carry out an attack is almost nonexistent. While script kiddies may not pose the same level of danger as organized criminals or a full nation state, they are still dangerous. In fact, professional hackers sometimes manipulate script kiddies into helping pull off their attacks (Alpine Security 2020).

According to Verizon's 2020 Data Breach Investigation Report, over half of all attacks are carried out by organized crime groups (Verizon 2020). Organized cyber criminals are responsible for incidents of all kinds - ransomware attacks, attacks that take down networks, information theft, identity theft, and financial crimes. Some use tooling they themselves built; others use publicly available or even stolen tools. Some cyber crime organizations even provide made to order malware as a service to other criminals (Fromiti 2020). Nation-states looking to expand their campaigns will also sometimes look to cyber crime organizations and sponsor them to act on their behalf (DoJ 2020).

Section IV: The Role of Offensive Capabilities

Regardless of how they're developed, offensive cyber capabilities of all varieties end up with relationships to nearly every other actor in the cyber threat actor network and are perhaps the most critical actor to understand. Consider the example of Mimikatz, a tool for stealing passwords, developed by hobbyist Benjamin Delpy (Greenberg 2017). Mimikatz may have been intended as a project for Delpy to familiarize himself with development and demonstrate a vulnerability in how computers stored passwords, but since then Delpy has acknowledged it's gone beyond that.

> *"Mimikatz wasn't at all designed for attackers. But it's helped them. When you create something like this for good, you know it can be used by the bad side too." - Benjamin Delpy*

Mimikatz is now integrated into a significant portion of tooling, and even tooling that doesn't use it directly uses variants. Even malware not using variants of Mimikatz sometimes still has its credential stealing components heavily influenced by the design of Mimikatz. Microsoft

ultimately redesigned how it handles credentials in response to Mimikatz, and networks around the world had to adjust their practices to ensure they were protected. This story is illustrative, but it is not unique, nor is this problem limited to tooling. As mentioned earlier, the GCHQ identified a critical flaw in a service used heavily throughout the world and reported it to the maintainer. In mid-May of 2019, Microsoft released an update to fix it (Msrc 2019). By late October, bad actors had teased out what the vulnerability was and developed the capability to exploit it (ESentire 2020). Had the vulnerability never been disclosed, it may never have even been discovered.

A recent study from the RAND corporation concludes that it may not always be in a nation-state's best interest to publicly disclose a vulnerability. As the leader of the study put it, "publicly disclosing a vulnerability that isn't known by one's adversaries gives them the upper hand, because the adversary could then protect against any attack using that vulnerability, while still keeping an inventory of vulnerabilities of which only it is aware of in reserve. In that case, stockpiling would be the best option." (Continuity Central 2020) Unfortunately, there's more to it than that. A collection of major vulnerabilities leading to widespread ransomware attacks was disclosed in 2017 by a group known as the Shadow Brokers, claiming it was stolen from the NSA (SentinelOne 2020). Regardless of whether their claim is true, it highlights a major drawback of stockpiling capabilities. If released suddenly, attackers will have an immediate advantage over the vulnerable networks. The same has happened with tooling; in 2020, FireEye was compromised by attackers that stole their custom tooling used for penetration testing and red team engagements (Bing & Menn 2020).

An interesting phenomenon arises from the network described in Section III, and it gives insight into some of the effects of research into offensive security and offensive cyber capabilities development. As companies' defensive posture improves, attackers are forced to

work to develop new tools, techniques, and procedures to breach their networks. As this happens, defenders must also adapt to the attackers' new techniques in order to remain safe. Without the ability to properly audit their network, companies can fall behind attackers and linger in the phase where attackers have the advantage. As described in Section III, penetration testing teams and red teams rely heavily on the offensive cyber capabilities available to them. The cost of developing custom tools is high in time, manpower, and maintenance. Thus, in most scenarios, these teams depend on publicly and commercially available tooling. Without this tooling available, the cycle of improvement in the actor network slows down, leaving the attackers with the advantage. Yet, it's not always that simple. With available tooling that represents the industry's best efforts to emulate the capabilities of an attacker, malicious actors are left with the perfect platform upon which they can build their next generation of capabilities. Thus, the tooling required to ensure defenders can keep pace with attackers can also serve to increase the rate at which attackers grow their own capabilities. This cycle makes halting the development of publicly and commercially available offensive cyber capabilities a dangerous task.

Section V: Conclusion

Actor Network Theory gives insight into the complex actor network involving people, organizations, and tooling that defines the cyber threat landscape. Understanding this actor network illuminates the numerous impacts of publicizing and commercializing offensive cyber capabilities. This actor network, like any, has incredible complexity which is impossible to fully capture in any paper, so while it is hoped the reader will find the insights valuable, this research does not present a complete picture. Future work may seek to examine the role of public opinion in this network and the effects of public offensive capabilities may have on public opinion.

WORKS CITED

Accenture Security. (2020). *2020 CYBER THREATSCAPE REPORT* (Rep.). Retrieved https://www.accenture.com/_acnmedia/PDF-137/Accenture-2020-Cyber-Threatscape-Report.pdf

Alpine Security. (2020, August 02). Do Script Kiddies Carry Out Most Cyber Attacks? Retrieved from https://alpinesecurity.com/blog/do-script-kiddies-carry-out-most-cyber-attacks/

Bing, C., & Menn, J. (2020, December 08). U.S. cybersecurity firm FireEye discloses breach, theft of hacking tools. Retrieved from https://www.reuters.com/article/us-fireeye-cyber/u-s-cybersecurity-firm-fireeye-discloses-breach-theft-of-hacking-tools-idUSKBN28I31E

CheckPoint. (2020, May 15). About Us. Retrieved from https://research.checkpoint.com/about-us/

Continuity Central. (2017, March 10). Study looks at zero-day vulnerabilities and what entities do when they discover them. Retrieved from https://www.continuitycentral.com/index.php/news/technology/1821-study-looks-at-zero-day-vulnerabilities-and-what-entities-do-when-they-discover-them

ESentire. (2020, January 11). UPDATE: BlueKeep Active Exploitation. Retrieved from https://www.esentire.com/security-advisories/bluekeep-active-exploitation-update

SentinelOne. (2020, May 04). Eternalblue: The NSA-developed Exploit That Just Won't Die. Retrieved from https://www.sentinelone.com/blog/eternalblue-nsa-developed-exploit-just-wont-die/

FireEye. (2021). Network Security Penetration Testing: Mandiant. Retrieved from https://www.fireeye.com/mandiant/penetration-testing.html

FireEye. (2021). Endpoint Security Software and Solutions. Retrieved from https://www.fireeye.com/products/endpoint-security.html

Fromiti. (2020, April). Organized Crime / Cybercrime Module 13 Key Issues: Cyber Organized Crime Activities. Retrieved from https://www.unodc.org/e4j/en/cybercrime/module-13/key-issues/cyber-organized-crime-activities.html

Geers, K., Kindlund, D., Moran, N., & Rachwald, R. (2014). World War C: Understanding Nation-State Motives Behind Today's Advanced Cyber Attacks. *FireEye White Paper*. Retrieved March 12, 2021, from https://www.fireeye.com/content/dam/fireeye-www/global/en/current-threats/pdfs/fireeye-wwc-report.pdf.

Greenberg, A. (2017, November 9). How the Mimikatz Hacker Tool Stole the World's Passwords. Retrieved from https://www.wired.com/story/how-mimikatz-became-go-to-hacker-tool/

HackerOne. (2021). Bug Bounty Programs For Businesses. Retrieved from https://www.hackerone.com/product/bounty

Hedström, K., Dhillon, G., & Karlsson, F. (2010). Using Actor Network Theory to Understand Information Security Management. *Security and Privacy – Silver Linings in the Cloud IFIP Advances in Information and Communication Technology,* 43-54. doi:10.1007/978-3-642-15257-3_5

Higgins, K. J. (2017, April 13). Nation-State Hackers Go Open Source. Retrieved from https://www.darkreading.com/threat-intelligence/nation-state-hackers-go-open-source/d/d-id/1328619

Kaspersky. (2021, January 13). What is WannaCry ransomware? Retrieved from https://usa.kaspersky.com/resource-center/threats/ransomware-wannacry

Kim, P. (2018). *The hacker playbook 3: Practical guide to penetration testing*. North Charleston, SC: Secure Planet LLC.

Knowles, S. (2020, November 17). What Is A Threat Actor?: Cyber Threat Actors: Blog. Retrieved from https://www.nexor.com/what-is-a-threat-actor/

Microsoft. (2021). Microsoft Security Vulnerability Research (MSVR). Retrieved from https://www.microsoft.com/en-us/msrc/msvr

Msrc. (2019, June 20). Microsoft Security Response Center. Retrieved from https://msrc-blog.microsoft.com/2019/05/14/prevent-a-worm-by-updating-remote-desktop-services-cve-2019-0708/

National Security Agency. (2015, June 5). Secure Host Baseline - National Security Agency - Applications. Retrieved from https://apps.nsa.gov/iaarchive/library/ia-guidance/security-tips/secure-host-baseline.cfm

National Security Agency. (2020, January 14). *Patch Critical Cryptographic Vulnerability in Microsoft Windows Clients and Servers* (Rep.). doi:https://media.defense.gov/2020/Jan/14/2002234275/-1/-1/0/CSA-WINDOWS-10-CRYPT-LIB-20190114.PDF

O'Donnell, A. L., & O'Donnell, L. (2021, January 4). Ticketmaster Coughs Up $10 Million Fine After Hacking Rival Business. Retrieved from https://threatpost.com/ticketmaster-10-million-fine-hacking-rival/162695/

O'Malley, M. (2020, March 26). Concerned about Nation State Cyberattacks? Here's how to Protect Your Organization. Retrieved from https://www.securitymagazine.com/articles/91889-concerned-about-nation-state-cyberattacks-heres-how-to-protect-your-organization

Project Zero. (2021). About Project Zero. Retrieved from https://googleprojectzero.blogspot.com/p/about-project-zero.html

Saminottawa. (2020, September 08). Inside the Chinese military attack on Nortel. Retrieved from https://globalnews.ca/news/7275588/inside-the-chinese-military-attack-on-nortel/

Defense Information Systems Agency. (2021, March 08). Security Technical Implementation Guides. Retrieved from https://public.cyber.mil/stigs/

Schectman, J., & Bing, C. (2019, January 30). Exclusive: UAE used cyber super-weapon to spy on iPhones of foes. Retrieved from https://www.reuters.com/article/us-usa-spying-karma-exclusive/exclusive-uae-used-cyber-super-weapon-to-spy-on-iphones-of-foes-idUSKCN1PO1AN

Tsing, W., & ABOUT THE AUTHOR William Tsing Breaking things and wrecking up the place since 2005. (2019, November 15). The Advanced Persistent Threat files: Lazarus Group. Retrieved from https://blog.malwarebytes.com/threat-analysis/2019/03/the-advanced-persistent-threat-files-lazarus-group/

US Department of Justice. (2020, February 13). Chinese Telecommunications Conglomerate Huawei and Subsidiaries Charged in Racketeering Conspiracy and Conspiracy to Steal Trade Secrets. Retrieved from https://www.justice.gov/opa/pr/chinese-telecommunications-conglomerate-huawei-and-subsidiaries-charged-racketeering

US Department of Justice. (2020, January 03). Countering State-Sponsored Cybercrime. Retrieved from https://www.justice.gov/usao-sdny/countering-state-sponsored-cybercrime

Verizon. (2020). *2020 Data Breach Investigations Report* (Rep.). Retrieved https://enterprise.verizon.com/resources/reports/2020-data-breach-investigations-report.pdf

Villarreal, R. (2019, January 16). Merlin The (C2) Wizard! Retrieved from https://bestestredteam.com/2019/01/16/merlin-the-c2-wizard/

Walsham, G. (1997). Actor-Network Theory and IS Research: Current Status and Future Prospects. *Information Systems and Qualitative Research,* 466-480. doi:10.1007/978-0-387-35309-8_23

National Cyber Security Centre. (2019, May 17). Weekly Threat Report 17th May 2019. Retrieved from https://www.ncsc.gov.uk/report/weekly-threat-report-17th-may-2019

Zorz, M., & 7, A. (2020, April 07). Full-time bug hunting: Pros and cons of an emerging career. Retrieved from https://www.helpnetsecurity.com/2020/04/07/bug-hunting-career/