Undergraduate Thesis Prospectus

Anomaly-Based Intrusion Detection for Linux Web Servers

(technical research project in Computer Science)

Harboring Malware for Good: Government Purchase of Zero-Days

(STS research project)

by

Roman Bohuk

October 31, 2019

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.

signed: _____ date: _____

approved: _____ date: _____
Peter Norton, Department of Engineering and Society

approved: _____ date: _____
Jack W. Davidson, Department of Computer Science

General research problem

How can cyber attacks be prevented?

Advancements in technology and proliferation of connected devices open up new attack surfaces and create new threats. According to the Internet Society, organizations have collectively lost over $45 billion due to breaches of confidentiality, integrity, and availability of their data in 2018. Despite awareness of the treat, cyberattacks are still the fastest growing crime in the United States, and losses are increasing (Summerville, 2017). In addition to local criminal groups, perpetrators include nation states, and attacks are not limited to purely virtual systems. Attacks on cyberphysical systems such as power grids, hospital buildings, and voting machines threaten human lives and national sovereignty. Cyber weapons have become indispensable to the military arsenals of various nation states, including the United States, and their power is comparable to that of physical weapons. The cybersecurity industry combats these threats, yet it is lagging behind. Cyber attacks threaten both individuals and groups, so proper preventative measures are needed.

Anomaly-based intrusion detection for Linux web servers

How can malicious activity be identified on Linux web servers through anomaly detection?

I will work under the supervision of Prof. Jack W. Davidson in the Computer Science department. This will be an independent project, but it aligns closely with my capstone class.

A web server is a piece of software that communicates with a client program over the internet through a series of requests. An attacker can manually craft a malicious request to perform an unintended action on the web server in an attempt to steal data or disrupt the service.

A study by Positive Technologies in 2018 found that 19 percent of tested web applications have vulnerabilities that enable an attacker take control of the server and the average number of vulnerabilities jumped dramatically from the year prior (Positive Technologies, 2018). These findings are staggering especially since the majority of those applications store and process user data. The number of web applications increases as internet usage grows and companies ramp up their infrastructure. This causes an explosion of traffic volume that the security teams have to analyze and process. Combined with a workforce shortage (in almost every position within cybersecurity), these cyber threats necessitate more automated solutions (Crumpler & Lewis, 2019).

One method to detect these exploits programmatically is to implement functions that inspect each request for signs of known malicious activity. This is known as rule-based detection, and it requires extensive signature databases that must stay up-to-date. Thus, it is powerless against new attacks and zero-day exploits. The alternative is anomaly detection. The system would observe benign traffic to form statistical models and raise flags if it sees something out of the ordinary.

The goal of this project is to move away from regular, rule-based detection methods that are common on the market (commonly referred to as intrusion detection systems and web application firewalls) and explore anomaly detection as an approach for web application security. The concept of anomaly detection in security has been used for at least 30 years, but most of the research has been done in general network security rather than web security specifically (Winkler & Page, 1989; Cho & Cha, 2004). In 2003, Kruegel and Vigna claimed to present the "first anomaly detection system specifically tailored to the detection of web-based attacks." Since then, multiple papers have been published focusing on various distinct parts of a web

request and increasingly utilizing machine learning models in more recent papers (Wen, Guo, & Yu, 2013; Zhang, et al., 2015; Zhang, Lu, & Xu, 2017). I plan to review and reproduce some of those methods. Also, a lot of previous research performs analysis on logs to identify malicious activity, but I hope to figure out a better, non-intrusive method for collecting requests on-the-fly. The end goal is to come up with a running and expandable tool that attempts web intrusion detection.

The OWASP Foundation maintains a comprehensive list of known attacks, and it will be a good reference for this project. I have access to a few production web servers that I can use for data. I also plan to set up various vulnerable honeypots to attract real traffic and collect requests from automated tools.

Harboring malware for good: government purchase of zero-days

How do governments rationalize buying critical security vulnerabilities (zero-days) and hiding them from affected companies?

Should any government disclose software vulnerabilities to vendors or retain and stockpile them (Ablon & Bogart, 2017)? If they are released, the affected company can patch the flaws immediately. Otherwise, attackers have the time to discover these flaws themselves and attack companies before the vendors can fix them.

This problem also has controversial implications. With a secret backdoor to a company's data, the government intelligence agencies can obtain personal information of its citizens. When FBI asked Apple to unlock an iPhone to help with a terrorist case, the company refused (Nakashima, 2016). Apple's CEO Tim Cook (2016) argued that a secret key to a phone would only be "as secure as the protections around it" and it would be equivalent to a master key

"capable of opening hundreds of millions of locks — from restaurants and banks to stores and homes." That key could be leaked, or the attackers could reverse engineer it for themselves. When a group called Shadow Brokers released a portion of NSA cyberweapon code on the internet, it became the standard offensive tool for exploiting Windows SMB service (Newman, 2018), and it is still in use today. Government purchase of zero-days gives these exploits a price tag and turns them into a valuable commodity. This incentivizes agencies to stockpile and develop these offensive capabilities instead of helping patch vulnerabilities, aligning governments' interests with those of the criminals.

Ablon & Bogart of the RAND Corporation (2017) propose that companies cannot rely on responsible disclosure and should instead invest in novel intrusion detection mechanisms. They argue that the merits of a zero-day release to the vendor depend on whether or not other malicious groups are aware of it, but this is difficult to gauge. Research shows that approximately 5.7 percent of any given set of vulnerabilities have been discovered by others after a year. Emery (2017) suggests that by purchasing zero-days, governments encourage responsible disclosure by private researchers motivated by money, which keeps the vulnerabilities off the black market. Stockton and Golabek-Goldman (2013) examine the "anarchic" black market for zero-day exploits and the government's role in it.

Zero-day exploits open secret backdoors to a target's data and threaten the integrity of its services. Companies offer researchers bug-bounty programs (HackerOne, 2017; Ahmed, 2015) to deter them from selling their findings to governments. Like other participant groups, they want these findings to go directly to them.

Some vulnerability researchers who identify flaws for a living may not care how their work is used (Barth, 2019). Others seek to reduce governments' access to spying tools by

identifying flaws before others do (Greenberg, 2014). Groups such as Google Project Zero fund and hire researchers to find and report vulnerabilities in commonly used tools regardless of whether or not they are owned by Alphabet Inc. (Google, 2019).

Resellers, escrows, and brokers facilitate legal and illegal transactions and protect parties' anonymity. Some maintain a private referral network of clients with entrance fees (Mitnick, 2019) while others help auction the exploit "equity" to the highest bidder (Anderson, 2007) for profit.

Government intelligence services are also participants (Zetter, 2014). They use zero-days to obtain intelligence, thwart crime, and attack other states. Like conventional weapons, exploits in cyber space can advance national agendas. Former White House cybersecurity coordinator Rob Joyce revealed the "tension between the government's need to sustain the means to pursue rogue actors in cyberspace through the use of cyber exploits, and its obligation to share its knowledge of flaws" to "ensure digital infrastructure is upgraded and made stronger in the face of growing cyber threats."

References

Ablon, L., & Bogart, A. (2017). *Zero Days, Thousands of Nights: The Life and Times of Zero-Day Vulnerabilities and Their Exploits*. Santa Monica, CA: Rand Corporation.

Ahmed, M. (2015, April 10). Internet companies pay out to those who spot bugs. *Financial Times*, https://www.ft.com/content/fcd027b4-c0d7-11e4-876d-00144feab7de

Anderson, N. (2007, July 9). WabiSabiLabi wants to be the eBay of 0-day exploits. *ArsTechnica.* https://arstechnica.com/information-technology/2007/07/wabisabilabi-wants-to-be-the-ebay-of-0-day-exploits/

Barth, B. (2019, August 8). Selling zero-days to governments takes some business savvy, says former bug broker. *SC Magazine*. https://www.scmagazine.com/home/security-news/vulnerabilities/selling-zero-days-to-governments-takes-some-business-savvy-says-former-bug-broker/

Cho, S., & Cha, S. (2004). SAD: web session anomaly detection based on parameter estimation. *Computers & Security*, 23(4), 312–319. doi: 10.1016/j.cose.2004.01.006

Cook, T. (2016, February 16). A Message to Our Customers. *Apple*. https://www.apple.com/customer-letter/

Crumpler, W., & Lewis, J. A. (2019, January 29). The cybersecurity workforce gap. *Center for Strategic and International Studies*. https://csis-prod.s3.amazonaws.com/s3fs-public/publication/190129_Crumpler_Cybersecurity_FINAL.pdf

Emery, A. C. (2017). Zero-day responsibility: the benefits of a safe harbor for cybersecurity research. *Jurimetrics, 57*(4), 483-503. https://search.proquest.com/docview/1965541181

Google (2019, July 31). Vulnerability Disclosure FAQ. *Project Zero*. https://googleprojectzero.blogspot.com/p/vulnerability-disclosure-faq.html.

Greenberg, A. (2014, July 15). Meet 'Project Zero,' Google's Secret Team of Bug-Hunting Hackers. *Wired*. https://www.wired.com/2014/07/google-project-zero/.

HackerOne (2017, January 3). Together We Hit Harder: HackerOne Company Values. *HackerOne Blog*, https://www.hackerone.com/blog/Together-We-Hit-Harder-HackerOne-Company-Values.

Internet Society. (2019). 2018 cyber incident & breach trends report. *Internet Society's Online Trust Alliance*. https://www.internetsociety.org/wp-content/uploads/2019/07/OTA-Incident-Breach-Trends-Report_2019.pdf

Joyce, R. (2017, December 10). Improving and making the vulnerability equities process transparent is the right thing to do. *White House.* https://www.whitehouse.gov/articles/improving-making-vulnerability-equities-process-transparent-right-thing/

Kruegel, C., & Vigna, G. (2003). Anomaly detection of web-based attacks. *Proceedings of the 10th ACM conference on Computer and communication security*. doi:10.1145/948143.948144

Menn, J. (2013, May 10). Special report: U.S. cyberwar strategy stokes fear of blowback. *Reuters.* https://www.reuters.com/article/us-usa-cyberweapons-specialreport/special-report-u-s-cyberwar-strategy-stokes-fear-of-blowback-idUSBRE9490EL20130510

Mitnick, K. (2019). Absolute Zero-Day Exploit Exchange | Premium Marketplace. *The Global Ghost Team*. https://www.mitnicksecurity.com/shopping/absolute-zero-day-exploit-exchange

Nakashima, E. (2016, February 17). Apple vows to resist FBI demand to crack iPhone linked to San Bernardino attacks. *Washington Post.* http://wapo.st/1TpBbL8

Newman, L. (2018, March 7). How leaked NSA spy tool 'EternalBlue' became a hacker favorite. *Wired.* https://www.wired.com/story/eternalblue-leaked-nsa-spy-tool-hacked-world/

OWASP Foundation. (2017). The ten most critical web application security risks. *The Open Web Application Security Project Foundation.* https://www.owasp.org/images/7/72/OWASP_Top_10-2017_%28en%29.pdf.pdf

Positive Technologies. (2019, March 5). Web application vulnerabilities: statistics for 2018. *Positive Technologies*. https://www.ptsecurity.com/ww-en/analytics/web-application-vulnerabilities-statistics-2019/

Shieh S., & Gligor V. (1992). Pattern-oriented intrusion-detection system and method. United States Patent US5278901A. https://patents.google.com/patent/US5278901A/en

Stockton, P. N., & Golabek-Goldman, M. (2013). Curbing the Market for Cyber Weapons. *Yale Law & Policy Review*, *32*(1), 239–266. https://digitalcommons.law.yale.edu/ylpr/vol32/iss1/11/

Summerville, A. (2017, July 25). Protect against the fastest-growing crime: cyber attacks. *CNBC.* https://www.cnbc.com/2017/07/25/stay-protected-from-the-uss-fastest-growing-crime-cyber-attacks.html

Wen, K., Guo, F., & Yu, M. (2013). Adaptive anomaly detection method of web-based attacks. *Journal of Computer Applications*, 32(7), 2003-2006. doi:10.3724/sp.j.1087.2012.02003

Winkler, J., & Page, W. (1989). Intrusion and anomaly detection in trusted systems. *[1989 Proceedings] Fifth Annual Computer Security Applications Conference*. doi:10.1109/csac.1989.81023

Zetter, K. (2014, November 17). U.S. gov insists it doesn't stockpile zero-day exploits to hack enemies. *Wired*. https://www.wired.com/2014/11/michael-daniel-no-zero-day-stockpile/

Zhang, M., Lu, S., & Xu, B. (2017). An Anomaly Detection Method Based on Multi-models to Detect Web Attacks. *2017 10th International Symposium on Computational Intelligence and Design (ISCID).* doi:10.1109/iscid.2017.223

Zhang, S., Li, B., Li, J., Zhang, M., & Chen, Y. (2015). A Novel anomaly detection approach for mitigating web-based attacks against clouds. *2015 IEEE 2nd International Conference on Cyber Security and Cloud Computing*. doi:10.1109/cscloud.2015.46