POSTER: A Multi-phased Multi-faceted IoT Honeypot Ecosystem

Armin Ziaie Tabari
University of South Florida
Tampa, Florida, USA
aziaietabari@usf.edu

Xinming Ou
University of South Florida
Tampa, Florida, USA
xou@usf.edu

ABSTRACT

The rapid growth of Internet of Things (IoT) devices makes it vitally important to understand real-world cybersecurity threats to them. Traditionally, honeypots have been used as decoys to mimic real devices on a network and help researchers/organizations understand the dynamic of threats. A crucial condition for a honeypot to yield useful insights is to let attackers believe they are real systems used by humans and organizations. However, IoT devices pose unique challenges in this respect, due to the large variety of device types and the physical-connectedness nature. In this work, we (1) presented an approach to create a multi-phased multi-faceted honeypot ecosystem, where researchers gradually increase the sophistication of a low-interaction IoT honeypot by observing real-world attackers' behaviors, (2) built a low-interaction honeypot for IoT cameras that allowed researchers to gain a concrete understanding of what attackers were going after on IoT camera devices, and (3) designed a proxy instance, called *ProxyPot*, that sits between IoT devices and the external network and helps researchers study the IoT devices' inbound/outbound communication. We used PorxyPot as a means to understanding attacks against IoT cameras and increasing the honeypot's sophistication. We deployed honeypots for more than two years. Our preliminary results showed that we were able to attract increasingly sophisticated attack data in each new phase. Moreover, we captured activities that appeared to involve direct human interactions rather than purely automated scripts.

CCS CONCEPTS

• **Security and privacy** → **Intrusion detection systems**; **Firewalls**; **Economics of security and privacy**.

KEYWORDS

Internet of Things; Honeypot

ACM Reference Format:
Armin Ziaie Tabari and Xinming Ou. 2020. POSTER: A Multi-phased Multi-faceted IoT Honeypot Ecosystem. In *Proceedings of the 2020 ACM SIGSAC Conference on Computer and Communications Security (CCS '20), November 9–13, 2020, Virtual Event, USA.* ACM, 3 pages. https://doi.org/10.1145/3372297.3420023

1 INTRODUCTION

Over the past few years, IoT devices have become an essential tool in people's daily activities. It was estimated that by 2025, there will be at least 41.6 billion IoT devices connected to the Internet [1], a 512% increase compared to 2018 (8 billion IoT devices) [2]. The exponential growth poses grave concerns regarding new security threats. Most IoT devices have simple accessible vulnerabilities including default username and password and open telnet/ssh port, to name just two. We are unfortunately at a time when exposure to attacks against IoT devices has become a reality, if not worse compared to traditional computing systems. Each new IoT device could offer a new passageway to adversaries and expose the entire network. For instance, more than 20% of companies around the world have experienced at least one IoT-related attack in the past few years [3, 4].

Creating an effectual cyber-security procedure or product needs a thorough understanding of the existing and possible threats. IoT has become an interesting new target for adversaries. It is thus highly crucial to understand what those attackers want, as well as their *modus operandi*. For a long time honeypots have helped security researchers to understand various types of attacker behaviors. By analyzing data captured by honeypots (network logs, downloaded files, *etc.*), researchers can uncover new tools and methods used by hackers, attack trends, and zero-day vulnerabilities. This information is highly valuable to improve cybersecurity measures, especially when organizations are resource-strapped when it comes to fixing security vulnerabilities.

There are two main challenges for creating IoT honeypots that can yield useful data for research:

(1) The types of different IoT devices are vast, each of which has unique features that an attacker may wish to access. It is infeasible to build one honeypot system that can capture even a significant portion of all IoT devices. Thus, we adopt a *multi-faceted* approach to IoT honeypot engineering.
(2) The specific nature of attackers' activities towards IoT devices is largely unknown at this point, and there could be very different focuses on the attacker's side. Moreover, richness of the response from an IoT device is much greater than traditional IT systems due to the interaction with the physical world. It would require significant amount of engineering work to emulate those different types of responses for different devices. Thus, we adopt a *multi-phased* approach where the sophistication of the emulated responses are gradually increased as data is gathered and analyzed to understand what the attackers might be going after.

In this work, we present our approach towards a comprehensive experimentation and engineering framework for capturing and analyzing real-world cyber-attacks on IoT devices using honeypots.

CCS '20, November 9–13, 2020, Virtual Event, USA
© 2020 Copyright held by the owner/author(s).
ACM ISBN 978-1-4503-7089-9/20/11.
https://doi.org/10.1145/3372297.3420023

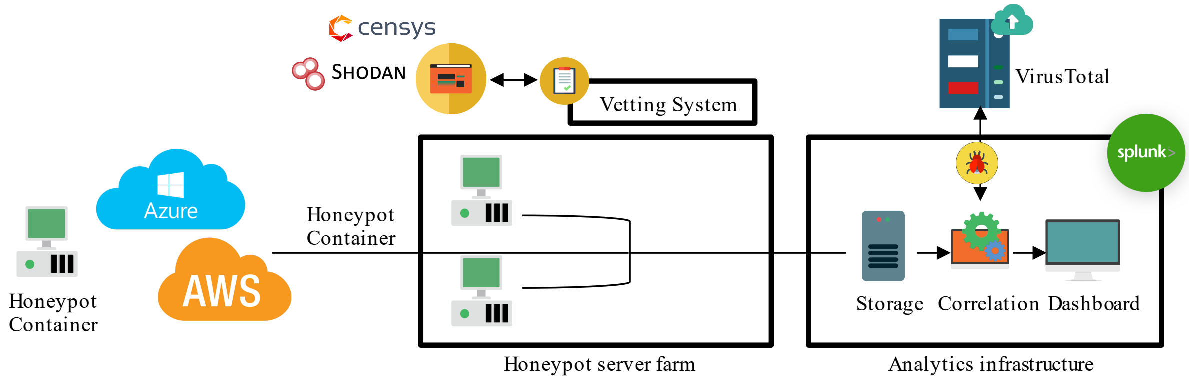

Figure 1: IoT Honeypot Ecosystem

2 IOT HONEYPOT ECOSYSTEM

To have a successful honeypot environment for research in IoT security, just having boxes running different emulated or real IoT devices is not sufficient. The honeypots need to be carefully maintained and monitored to allow intelligent adaptation on the way they respond to different types of traffic, so an attacker can be "hooked" and made interested in further exploring it. Our goal was to create a carefully designed ecosystem where a variety of honeypot devices working together with a "vetting" and "analysis infrastructure" by their side. Figure 1 illustrates our implemented ecosystem, which consists of three separate components.

1) Honeypot server farms: In this study, we used both an on-premise server and cloud infrastructures. Using cloud servers helped us gain vantage points in different countries and cover a wide range of geographic locations.

2) Vetting system: A honeypot is valuable only as long as it remains undetectable. The vetting system ensures servers are sufficiently difficult for an adversary to detect them as a honeypot. For this purpose, we tried a number of fingerprinting approaches to vet any new honeypot instance. Both manual and automatic vetting were used (*e.g.*, through Metasploit). In addition, we used Internet device search engines Shodan [5] and Censys [6] to make sure they look like the real ones they imitate. Most importantly, we analyzed our honeypot logs to identify fingerprinting attempts. We then adapt our honeypots to render those fingerprinting approaches ineffective.

3) Analysis infrastructure: The success of a honeypot depends on two factors: 1) the way the honeypot software is developed and implemented; and 2) the log analysis process. Carefully analyzing the logs is as important as the honeypot development and implementation. We utilized Splunk [7] for log management and analysis. All the logs captured in our ecosystem are sent to our centralized Splunk server. We designed an app on Splunk that automated the analysis processes and used VirousTotal, DShield, and AbuseIPDB [1] to analyze attacker IP addresses and captured malicious files.

3 MULTI-FACETED AND MULTI-PHASED DEPLOYMENT

As explained in Section 1, given that different IoT devices have different specifications and configurations, each honeypot needs to be designed and configured in a unique way. Thus we adopted a multi-faceted approach to building the various honeypot instances. We also employ a multi-phased approach, where deployed honeypots' sophistication is gradually increased based on the observation of collected data.

In the first phase, we simply deployed the honeypots and collected data. In the second phase, we analyzed the captured data and tried to understand what information cybercriminals were looking for, to provide responses accordingly for the purpose of eliciting further adversary behaviors. These two phases iterate until we are satisfied with the insights gained. In the third phase, we used all the information we learned previously and deployed more advanced honeypots. This multi-phased investigation dovetails with our multi-faceted honeypot approach. We have so far designed three facets in our ecosystem. We both used off-the-shelf honeypot emulators and adapted them, and built specific emulators from scratch. The first facet we designed was *HoneyShell*. We utilize off-the-shelf Cowrie honeypots to emulate vulnerable IoT devices with open ssh (port 22) or telnet (port 23). In each phase, we gradually increased the sophistication of login credentials by analyzing previous phases' logs. By doing so, we were able to obtain increased sophistication and volume of commands executed in our honeypots. This provided us more insights into attackers compared to simply deploying a honeypot without adapting it, at the same time avoided wasting engineering effort in creating a more realistic emulation if the engineered features ended up not being explored by any attackers. *HoneyWindowsBox* was the second facet we deployed in our ecosystem to emulate IoT devices based on the Windows platform. Dionaea and KFSensor were chosen for this purpose. For our third facet, we developed a more specific honeypot from scratch to emulate behaviors of an IoT camera device. First, we developed *ProxyPot* (Figure 3), a proxy instance that sits between an IoT device and the network gateway and captures all traffic that goes between. Based on the insights gained through the ProxyPot data, we built a honeypot for IoT camera and coined it *HoneyCamera* (Figure 2). Honeycamera is a low-interaction honeypot for D-Link

[1] https://www.virustotal.com/, https://www.dshield.org/, and https://www.abuseipdb.com/ respectively.

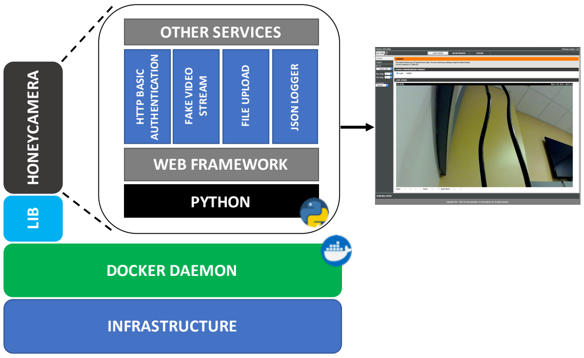

Figure 2: HoneyCamera architecture

IoT cameras. In *Phase 1*, three honeypots were deployed. Two of them were used to emulate D-Link DCS-5020L and the other one to imitate D-Link DCS-5030L camera. In *Phase 2*, based on the logs captured in phase 1, multiple two-step attacks were designed for the HoneyCamera. After that, we did see an attack that first exploited a well-known vulnerability that we planted inside the HoneyCamera, to read username and password. Then the same attacking IP logged in through the Camera login page using the stolen credential. More details on the HoneyCamera can be found in the accompanying technical report [8].

4 RELATED WORK

Luo *et al.* [9] designed an “intelligent-interaction” honeypot for IoT devices. It actively scanned other IoT devices around the world and sent part of received attacks to them as a means to eliciting legitimate responses. Such experiments need to be done with great care and ethical concerns, to prevent those devices from becoming unwitting victims. In our honeypot ecosystem, all attack traffic is forwarded to devices under our possession. The ProxyPot allows for controlled exposure of our own devices to attackers.

Wang *et al.* [10] presented a hybrid IoT honeypot framework called IoTCMal. It included a low-interaction component with telnet/ssh service and high-interaction vulnerable IoT devices and was used to collect and analyze malware samples. Feng *et al.* [11] used a honeypot system consisting of both real devices and simulated services in their IoTShield framework. The honeypot system collected real-world attack traffic in about two months’ time and the data was used to create attack signatures for automated protection. Vetterl *et al.* [12] used firmware images to emulate CPE/IoT devices and ran them as honeypots. Moreover, there have been a number of honeypot studies that utilized low-interaction and high-interaction honeypots separately or together to study attacks on IoT devices [13–18]. Compared to the prior works, our main contribution is the design, implementation, and deployment of a multi-phased multi-faceted honeypot ecosystem that addresses the challenges of capturing useful attack data on IoT devices, and studies adversaries’ behaviors as they evolve.

5 CONCLUSION

In this work, we present a multi-faceted and multi-phased approach to building an IoT honeypot ecosystem. An evolving honeypot ecosystem can attract more interesting attacks that could yield higher utility for research and operation, compared to one that is built once and deployed for a short period of time (*e.g.*, a couple of months). In particular, our approach seems to be uniquely capable of capturing human attack activities, as opposed to simply automated attack scripts.

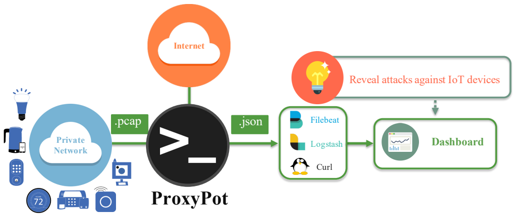

Figure 3: ProxyPot architecture

ACKNOWLEDGMENTS

We are grateful to Raj Rajagopalan for insightful comments and discussions throughout this work. We also thank Danielle Ward and Nathan Schurr for helpful input to this research.

REFERENCES

[1] “The growth in connected IoT devices is expected to generate 79.4ZB of data in 2025, according to a new IDC forecast,” Jun 2019.
[2] P. Newman, “The Internet of Things 2020: Here’s what over 400 IoT decision-makers say about the future of enterprise connectivity and how IoT companies can use it to grow revenue,” Mar 2020.
[3] “Cisco visual networking index: Forecast and trends, 2017-2022 white paper,” Feb 2019.
[4] “IoT heading for mass adoption by 2019,” Feb 2017.
[5] “The search engine for the internet of things,” https://www.shodan.io/.
[6] Z. Durumeric, D. Adrian, A. Mirian, M. Bailey, and J. A. Halderman, “A search engine backed by Internet-wide scanning,” in *the 22nd ACM Conference on Computer and Communications Security*, Oct. 2015.
[7] “SIEM, AIOps, application management, log management, machine learning, and compliance,” https://www.splunk.com/.
[8] A. Ziaie Tabari and X. Ou, “A first step towards understanding real-world attacks on IoT devices,” *arXiv e-prints*, Mar. 2020, https://arxiv.org/abs/2003.01218.
[9] T. Luo, Z. Xu, X. Jin, Y. Jia, and X. Ouyang, “IoTCandyJar: Towards an intelligent-interaction honeypot for IoT devices,” in *Black Hat USA*, 2017.
[10] B. Wang, Y. Dou, Y. Sang, Y. Zhang, and J. Huang, “IoTCMal: Towards a hybrid IoT honeypot for capturing and analyzing malware,” in *the IEEE International Conference on Communications (ICC)*, 2020.
[11] X. Feng, X. Liao, X. Wang, H. Wang, Q. Li, K. Yang, H. Zhu, and L. Sun, “Understanding and securing device vulnerabilities through automated bug report analysis,” in *the 28th USENIX Conference on Security Symposium*, 2019.
[12] A. Vetterl and R. Clayton, “Honware: A virtual honeypot framework for capturing CPE and IoT zero days,” in *the APWG Symposium on Electronic Crime Research (eCrime)*, 2019.
[13] Y. P. Minn, S. Suzuki, K. Yoshioka, T. Matsumoto, and C. Rossow, “IoTPOT: Analysing the rise of IoT compromises,” in *the 9th USENIX Workshop on Offensive Technologies (WOOT)*, 2015.
[14] J. D. Guarnizo, A. Tambe, S. S. Bhunia, M. Ochoa, N. O. Tippenhauer, A. Shabtai, and Y. Elovici, “Siphon: Towards scalable high-interaction physical honeypots,” in *the 3rd ACM Workshop on Cyber-Physical System Security*, 2017.
[15] S. Dowling, M. Schukat, and H. Melvin, “A zigbee honeypot to assess IoT cyberattack behaviour,” in *the 28th Irish Signals and Systems Conference (ISSC)*, 2017.
[16] S. Chamotra, R. K. Sehgal, S. Ror *et al.*, “Honeypot deployment in broadband networks,” in *the International Conference on Information Systems Security*, 2016.
[17] M. Wang, J. Santillan, and F. Kuipers, “Thingpot: An interactive Internet-of-Things honeypot,” *arXiv preprint arXiv:1807.04114*, 2018.
[18] M. A. Hakim, H. Aksu, A. S. Uluagac, and K. Akkaya, “U-pot: A honeypot framework for UPnP-based IoT devices,” in *the IEEE 37th International Performance Computing and Communications Conference (IPCCC)*, 2018.