YourThings: A Comprehensive Annotated Dataset of Network Traffic from Deployed Home-based IoT Devices

Omar Alrawi
alrawi@gatech.edu
Georgia Institute of Technology
Atlanta, GA, USA

Aaron Faulkenberry
afaulken@gatech.edu
Georgia Institute of Technology
Atlanta, GA, USA

Fabian Monrose
fabian@ece.gatech.edu
Georgia Institute of Technology
Atlanta, GA, USA

Manos Antonakakis
manos@gatech.edu
Georgia Institute of Technology
Atlanta, GA, USA

ABSTRACT

We describe a novel and public dataset for evaluating the security and privacy of home-based IoT devices known as YourThings [1]. The dataset is the result of thousands of research hours to build a large testbed with diverse home-based IoT devices. Since publishing the manuscript and dataset in the proceedings of the IEEE Security & Privacy 2019 (Oakland), our work has been cited over 300 times, used in over 20 research projects, incorporated by consumer technology advocacy groups (Wirecutter), and influenced standards at the National Institute of Standards and Technology (NIST) and the Internet Engineering Task Force (IETF). Beyond academic works, our dataset has been used by practitioners in industry, including McAfee, Hamilton Beach, and Aura.

ACM Reference Format:
Omar Alrawi, Aaron Faulkenberry, Fabian Monrose, and Manos Antonakakis. 2022. YourThings: A Comprehensive Annotated Dataset of Network Traffic from Deployed Home-based IoT Devices. In *Annual Computer Security Applications Conference (ACSAC '22), December 5–9, 2022, Austin, TX, USA.* ACM, New York, NY, USA, 4 pages. https://doi.org/10.1145/nnnnnnn.nnnnnnn

1 INTRODUCTION

Insecure home-based IoT devices have contributed to record-breaking attacks on critical internet infrastructure [1]. To begin to understand the underlying security vulnerabilities that manifest in home-based IoT devices, a relatively large-scale study of diverse home-based IoT devices must be carried out. Prior to our work, researchers conducted security evaluations on small sets of similar IoT devices in an ad hoc fashion [2]. More often than not, these works would not provide their dataset publicly for other researchers to use or extend. Even if those datasets were made available, they were not comprehensive enough to provide a comparative view across home-based IoT vendors and device types.

We addressed this gap by building a state-of-the-art testbed of off-the-shelf home-based IoT devices and evaluated their security. In doing so, we captured the entire evaluation process by recording the full-packet captures from 45 diverse home-based IoT devices. We overcame several challenges in building and deploying this testbed, including dealing with power, networking, cooling, and special requirements for different types of home-based IoT devices. We thoroughly documented our methodology, security scoring process, and device configurations, which we made public to researchers and practitioners. Since then, YourThings has garnered the attention of the research community, consumer technology advocacy groups, acclaimed media, and practitioners from the Internet Engineering Task Force (IETF).

Our paper documenting this experiment appears in the proceeding of the 2019 IEEE Security & Privacy, a tier one security conference as a systematization of knowledge. Since its publication, our paper and dataset have accumulated over 300 citations and used in over 20 projects by the academic community. The dataset of 150GB of full-packet network capture traces has been downloaded over 130 times. Consumer technology advocacy groups, such as the Wirecutter, have used our data and methodology to inform consumers about secure IoT devices and provide recommendations. Our work has been featured in acclaimed media such as Newsweek and the New York Times, which disseminate technical material to wider audiences. Since our initial experiments, we have grown our lab to over 100 devices, continued to evaluate the security of each device weekly and are in the process of releasing a longitudinal dataset to further help improve the security state of the home-based IoT ecosystem.

2 DATASET

The YourThings project conducted a comprehensive security evaluation of 45 diverse home-based IoT devices. The security evaluation is an internet protocol (IP)-based, which allowed us to record the entire process in full-packet capture (PCAP) traces. The security evaluation includes network vulnerability scans, network interception assessment, and device idle and active network profiles. The YourThings dataset consists of 13 PCAP files ranging between 10GB to 13GB in size for a total of approximately 150GB. The YourThings project published these PCAP traces publicly and provided supplement material to allow others to reproduce them. The supplement

[1] https://yourthings.info

ACSAC '22, December 5–9, 2022, Austin, TX, USA

ACM ISBN 978-x-xxxx-xxxx-x/YY/MM.
https://doi.org/10.1145/nnnnnnn.nnnnnnn

materials include device to IP address mapping, dates for each evaluation, security evaluation rubric, and scoring methodology.

2.1 Testbed

Our network setup has three main components, the IoT subnet, custom Linux gateway, and an assessment machine. The assessment machine runs all our evaluation tools and sits on the same subnet as the IoT devices. Our gateway is a Debian Jessie Linux machine, which manages the network services (DHCP, DNS, etc.) and connects the IoT subnet to the Internet. Additionally, our gateway full-packet captures all IP traffic originating from the IoT subnet. We used a 24-port switch to connect wired IoT devices via Ethernet and a wireless access point for devices that require 802.11 WLAN. All the IoT devices are assigned a static IP based on their MAC address. We aimed to use open-source tools for the security evaluation so that researchers have access to the necessary tools for reproduciblity. Unfortunately, we could not release our tools since they are experimental and built for our testbed, and we do not have the resources to provide support.

2.2 Tools

We used Nessus Scanner [3] to scan devices and cloud endpoints for service discovery, service profiling, and vulnerability assessment. Nessus Scanner annotates the CVE [4] information with the versions of running services and provides a summary of their security state. Nessus Scanner uses the CVSS [5] scoring system to rate the severity of the discovered vulnerability on a scale from one to ten and categorizes them into low, medium, high, and critical. We used MobSF [6], Qark [7], and services from Kryptowire [8] to statically and dynamically evaluate each mobile application for the IoT devices. We analyzed both, the Android and the iOS applications. We used Nessus Network Monitor [3], ntop-ng [9], Wireshark [10], and sslsplit [11] to profile the communication edges for each device. We manually inspected traffic and tested them for MITM attack using sslsplit.

2.3 Collection

To generate this dataset, we invested substantial engineering resources over the span of 18 months to build the testbed, the monitoring tools, and the automation of the security evaluation. The testbed lab contains real-world off-the-shelf devices that span six categories, namely appliances, cameras, home assistants, home automation, media, and network devices. Our testbed required space, power, cooling, and network provisioning. Additionally, we faced several challenges deploying diverse types of devices, such as vacuums, thermostats, and light bulbs. Each device type required special conditions to properly operate. For example, the smart Nest thermostat required a magnetic relay to emulate the presence of an HVAC system. The light bulbs required light bulb sockets to power each light bulb, which became difficult to manage as we scaled the lab. Smart vacuums require floor space to roam around and generate observable network activities. As we scaled our testbed, we faced power constraints (lack of outlets and amperage) and heating issues (poorly ventilated small office space) that we had to solve.

We manually deployed each device and recorded the configuration parameters, internet connectivity settings, and firmware updates. As for the data collection (monitoring), we automated the traffic capture at the egress point of the network (LAN-to-WAN) on a custom-built Linux router. For the active profiles, we manually exercised core functional features on each device and recorded their time. The collection effort was automated, but we labeled the traffic manually by observing idle/active states. Finally, for the security evaluation, we deployed automated network vulnerability scans for the devices and their cloud endpoints. Moreover, we conducted manual network interception between the networked components in an IoT deployment (device to cloud, device to mobile, and mobile to cloud) to test their susceptibility to man-in-the-middle (MiTM) attacks. For each interception assessment, we evaluated the certificate pinning configuration and protocol integrity.

2.4 Documentation

We documented the dataset, evaluation artifacts, and security scoring methodology. We published the documentation on the YourThings website and included examples to guide researchers in reproducing and extending our efforts. The dataset is annotated with date ranges for identifying when the device and network evaluation was completed. Each device configuration and network IP address are mapped in a machine-readable format (CSV) to enable automation. In addition, we provide the intermediary results for each device evaluation, including the raw artifacts extracted from the security assessment that is used to quantify the security score for the devices. Lastly, we provide a detailed rubric scoring system and show how to modify the rubric weights for different threat models, including off-path, on-path, and geographically near attacker.

3 IMPACT

The YourThings dataset has been downloaded over 130 times, which accounts for 20TB (130*150GB) of PCAP traces shared with the public. Over the past four years, our work has been cited over 300 times, used by more than 20 research projects, used as a reference for standards and protocol improvements (NIST and IETF), and incorporated by the Wirecutter review website [12]. In addition, our work has made headlines in Newsweek [13] and The New York Times [14].

Table 1: Citation sources by societies.

Publication Societies	Citations
IEEE	67
ACM	29
Springer	15
Elsevier	10
Usenix	9
NDSS	2
Others	172
Total	304

3.1 Academic Research

To highlight the impact of the citations beyond the numerical count, Table 1 shows the publication societies (conferences and journals) for the works that cite YourThings . We can see that IEEE society is

the highest, followed by ACM, Springer, and Elsevier. These publication societies are among the most active in the field of computer engineering and computer science. Additionally, eight book series have referenced our work. Our work has had a broad impact on the research community with citations from areas in privacy, system design, network security, energy, human-computer interaction, measurements, access control, digital health, artificial intelligence, cyber forensics, and software testing to name a few. The "Other" category includes dissertations, technical reports, Arxiv papers, and miscellaneous digital publications citations. It is clear from Table 1 that the impact of our work on academic research is substantial.

Table 2: A sample of 15 academic research projects using the YourThings dataset.

Paper	Year	Application
HomeSnitch [15]	2019	Identifying and Controlling IoT Behavior
Storming the Kasa [16]	2019	Reproducibility of Security Evaluation
Hestia [17]	2019	Defining Least Privilage Network Policy
Ask the Experts [18]	2020	Security and Privcy Device Labels
PingPong [19]	2020	Automated Device Signature Identification
ML Traffic Classification [20]	2020	IoT Traffic Classification
Automated Standards [21]	2020	Automation of Security Assessment
IoTFinder [22]	2020	DNS-based Device Identification
IoT ETEI [23]	2021	Device Identification
FIAT [24]	2021	Improved IoT Authentication System
ByteIoT [25]	2021	IoT Device Identification
Standards and Technology [26]	2021	Survey on IoT Security Evaluation
Survey on Device Behavior [27]	2021	Survey on IoT Device Behavior Identification
PinBall [28]	2021	IoT Device Event Identification

3.2 Applied Research

To provide more concrete examples of how our work has impacted research, Table 2 presents a sample of 15 works that incorporate our dataset into their research. Specifically, these research projects span different applications such as device identification, behavior detection and classification, security and privacy device labels, IoT authentication techniques, device behavior transparency, and surveys on security assessments and device identification. These publications appear in various top security venues like IEEE S&P (Oakland) and NDSS. Beyond academic works, our dataset has been used by practitioners in the industry. McAfee detection engineers used our dataset to prototype certificate signatures to reduce false positives when identifying IoT device traffic. We worked with a McAfee security engineer to answer questions about our dataset and how they can extract certificates associated with each device's communication. Hamilton Beach, an appliance manufacturer, leveraged our systematic security evaluation methodology to vet cloud platforms they intended to use with their smart-home connected products. Lastly, our dataset is helping researchers define a protocol for sending DNS messages over the Constrained Application Protocol [29].

3.3 Improving Consumer Security

The Wirecutter has incorporated our methodology and evaluation results to recommend secure light bulbs to consumers [12]. We worked with Hamilton Beach, a consumer appliance company, to apply our security assessment methodology and identify secure IoT deployment practices. We worked with Aura, a digital identity management company to recommend techniques for identifying and securing deployed home-based IoT devices running inside their customers' networks.

3.4 News Outlets

Our work has been featured in a Newsweek piece about the security and privacy of home-based IoT devices [13]. The article highlights the evolution of home-based IoT products and the risk associated with using these internet-connected devices. The New York Times featured our work on Amazon's Sidewalk project that uses nearby Amazon home-based IoT devices to share internet connections [14]. The article describes the new feature, informs users how to turn it off, and highlights our work to provide a security and privacy perspective. The Verge featured our work in reference to the privacy issues found in the Anker Eufy Security Camera. Eufy claimed the video feed is end-to-end encrypted (E2EE); however, that was not what the case [30].

4 CONCLUSION

Our large-scale security evaluation of diverse home-based IoT deployments has had a noticeable impact on academic research and the security of home-based IoT devices. This is enabled by the release of our dataset that has allowed others to reproduce and apply it to various applications (see Table 2). Since the dataset has been made public, we have served 130 downloads accounting for 20TB of shared data with the community. We plan to release additional longitudinal data to the research community to further security and privacy research on home-based IoT devices.

REFERENCES

[1] M. Antonakakis, T. April, M. Bailey, *et al.*, "Understanding the mirai botnet," in *26th USENIX security symposium (USENIX Security 17)*, 2017, pp. 1093–1110.

[2] O. Alrawi, C. Lever, M. Antonakakis, and F. Monrose, "Sok: Security evaluation of home-based iot deployments," in *2019 IEEE symposium on security and privacy (sp)*, IEEE, 2019, pp. 1362–1380.

[3] tenable, *Nessus Professional*, http://info.tenable.com/rs/934-XQB-568/images/NessusPro__DS__EN_v8.pdf, 2005.

[4] MITRE, *About CVE*, http://cve.mitre.org/about/index.html, 1999.

[5] FIRST, *Common Vulnerability Scoring System SIG*, https://www.first.org/cvss/, 2005.

[6] A. Abraham, *Mobile Security Framework (MobSF)*, https://github.com/MobSF/Mobile-Security-Framework-MobSF/blob/master/README.md, 2016.

[7] Linkedin, *QARK - Quick Android Review Kit*, https://github.com/linkedin/qark/blob/master/README.md, 2016.

[8] *Kryptowire EMM+S*, http://www.kryptowire.com/enterprise.php, 2011.

[9] ntop, *High-Speed Web-based Traffic Analysis and Flow Collection*, https://www.ntop.org/products/traffic-analysis/ntop/, 1998.

[10] G. Combs, *About Wireshark*, https://www.wireshark.org, 1998.

[11] D. Roethlisberger, *SSLsplit - transparent SSL/TLS interception*, https://www.roe.ch/SSLsplit, 2009.

[12] *How wirecutter vets the security and privacy of smart home devices*, Sep. 2020. [Online]. Available: %5Curl%7Bhttps://www.nytimes.com/wirecutter/blog/smart-home-security-privacy/%7D.

[13] A. Piore, *We're surrounded by billions of internet-connected devices. can we trust them?* Oct. 2019. [Online]. Available: %5Curl%7Bhttps://www.newsweek.com/2019/11/01/trust-internet-things-hacks-vulnerabilities-1467540.html%7D.

[14] *Amazon sidewalk will share your internet with strangers. it's not as scary as it sounds.* Jun. 2021. [Online]. Available: %5Curl%7Bhttps://www.nytimes.com/wirecutter/blog/amazon-sidewalk-review/%7D.

[15] T. OConnor, R. Mohamed, M. Miettinen, W. Enck, B. Reaves, and A.-R. Sadeghi, "Homesnitch: Behavior transparency and control for smart home iot devices," in *Proceedings of the 12th conference on security and privacy in wireless and mobile networks*, 2019, pp. 128–138.

[16] A. Halterman, "Storming the kasa? security analysis of tp-link kasa smart home devices," *Creat. Compon*, 2019.

[17] S. Goutam, W. Enck, and B. Reaves, "Hestia: Simple least privilege network policies for smart homes," in *Proceedings of the 12th Conference on Security and Privacy in Wireless and Mobile Networks*, 2019, pp. 215–220.

[18] P. Emami-Naeini, Y. Agarwal, L. F. Cranor, and H. Hibshi, "Ask the experts: What should be on an iot privacy and security label?" In *2020 IEEE Symposium on Security and Privacy (SP)*, IEEE, 2020, pp. 447–464.

[19] R. Trimananda, J. Varmarken, A. Markopoulou, and B. Demsky, "Packet-level signatures for smart home devices," in *Network and Distributed Systems Security (NDSS) Symposium*, vol. 2020, 2020.

[20] V. Melnyk, P. Haleta, and N. Golphamid, "Machine learning based network traffic classification approach for internet of things devices," *Theoretical and Applied Cybersecurity*, vol. 2, no. 1, 2020.

[21] A. N. d. P. T. Gurgo *et al.*, "Automated standard based security assessment for iot," 2020.

[22] R. Perdisci, T. Papastergiou, O. Alrawi, and M. Antonakakis, "Iotfinder: Efficient large-scale identification of iot devices via passive dns traffic analysis," in *2020 IEEE European Symposium on Security and Privacy (EuroS&P)*, IEEE Computer Society, 2020, pp. 474–489.

[23] F. Yin, L. Yang, Y. Wang, and J. Dai, "Iot etei: End-to-end iot device identification method," in *2021 IEEE Conference on Dependable and Secure Computing (DSC)*, IEEE, 2021, pp. 1–8.

[24] Y. Xiao and M. Varvello, "Fiat: Frictionless authentication of iot traffic," in *Proceedings of the 17th International Conference on emerging Networking EXperiments and Technologies*, 2021, pp. 483–484.

[25] C. Duan, H. Gao, G. Song, J. Yang, and Z. Wang, "Byteiot: A practical iot device identification system based on packet length distribution," *IEEE Transactions on Network and Service Management*, 2021.

[26] B. Delinchant and J. Ferrari, "Standards and technologies from building sector, iot, and open-source trends," in *Towards Energy Smart Homes*, Springer, 2021, pp. 49–111.

[27] P. M. S. Sánchez, J. M. J. Valero, A. H. Celdrán, G. Bovet, M. G. Pérez, and G. M. Pérez, "A survey on device behavior fingerprinting: Data sources, techniques, application scenarios, and datasets," *IEEE Communications Surveys & Tutorials*, vol. 23, no. 2, pp. 1048–1077, 2021.

[28] C. Duan, S. Zhang, J. Yang, Z. Wang, Y. Yang, and J. Li, "Pinball: Universal and robust signature extraction for smart home devices," in *2021 IFIP/IEEE International Symposium on Integrated Network Management (IM)*, IEEE, 2021, pp. 1–9.

[29] M. S. Lenders, C. Amsüss, C. Gündoğan, T. C. Schmidt, and M. Wählisch, "DNS over CoAP (DoC)," RFC Editor, RFC Draft, Jul. 2022. [Online]. Available: %5Curl%7Bhttps://datatracker.ietf.org/doc/draft-lenders-dns-over-coap/%7D.

[30] *Anker's eufy lied to us about the security of its security cameras*, Dec. 2022. [Online]. Available: %5Curl%7Bhttps://www.theverge.com/2022/11/30/23486753/anker-eufy-security-camera-cloud-private-encryption-authentication-storage%7D.