Network Security for Home IoT Devices Must Involve the User: a Position Paper

Lorenzo De Carli and Antonio Mignano

Worcester Polytechnic Institute, Worcester MA 01609, USA

Abstract. Many home IoT devices suffer from poor security design and confusing interfaces, lowering the bar for successful cyberattacks. A popular approach to identify compromised IoT devices is network-based detection, in which network traffic is analyzed to fingerprint and identify such devices. However, while several network-based techniques for identifying misbehaving devices have been proposed, the role of the user in remediating IoT security incidents has been conspicuously overlooked. In this paper, we argue that successful IoT security must involve the user, even if the user is not a technical expert, and that the form in which security findings are communicated is as important as the technique used to generate such warnings. Finally, we present the design of a research testbed designed to foster further research in IoT security warnings.

1 Introduction

Increasingly, modern homes are incorporating internet-connected sensors and actuators such as smart cameras, lights and locks. These devices promise to bring a host of benefits, such as simplifying management of the home. Unfortunately, the potential benefits of IoT devices are increasingly outweighed by cybersecurity risks. Many IoT devices suffer from poor security design that render them vulnerable to cyberattacks [8, 20]. Furthermore, the complexity of a residential network with even just a handful of devices can quickly outgrow the network security skills of a non-technical owner.

Existing efforts to secure the smart home include IoT intrusion detection systems (IDS) which flag abnormal device behavior that may be symptom of compromise, based on analysis of device network traffic(e.g., [17,19]). Typically, in presence of a network attack the IDS is expected to generate a warning signal, but existing works tend to sidestep the question of *who* should receive these warnings and act upon them. Interesting, the intrusion detection techniques now proposed for the IoT network security space (attack signatures, network anomaly detection) tend to be adaptations of approaches previously originated in the enterprise space. However, enterprise network security is based on a specific set of assumptions, which do not necessarily hold true in home IoT security.

One of such assumptions is particularly problematic: the availability of an expert user (network administrator, security officer, etc.) who can understand and review the output of an IDS, and take action if necessary. The fact that such

an expert *does not* exist in most home setups raises a number of issues, which must be solved to make intrusion detection practical in the home IoT space.

The first issue is what to do with warnings generated by an IDS. The option of letting the system autonomously dealing with attacks (e.g., by blocking malicious flows) is tempting, but we argue that is not viable, both due to infeasibility (no IDS is correct all the time), and for incompatibility with what we consider an important design tenet, that a user must be able to retain *control* of their devices (Section 3). Having argued the necessity to keep the human in the loop, we then focus on a second issue: how to precisely involve the users in management of security incidents. Through example scenarios, we highlight that determining how to bring the user in the security loop is a hard problem which warrants further research (Section 4). Furthermore, in order to foster further research in IoT network monitoring systems with human in the loop, we describe a testbed for user-centric research on IoT security notifications and interactions (Section 5). Finally, using this testbed we collect preliminary user impressions (Section 6).

2 Background

In this section, we introduce the threat model for home IoT networks which is assumed throughout the rest of this paper, and we review related work in the area of network security.

2.1 Assumptions and Threat Model

Our goal is to provide guidance to designers of systems that aim at detecting compromised IoT devices via analysis of their traffic. We assume an attacker that can send arbitrary traffic to the home network, with the goal of acquire partial or full control of one or more devices. We assume the detection system has visibility over the traffic within the home network. The system may detect attempts to compromise devices, compromised devices, or both. It would typically run on the residential gateway, but it may also be offloaded to the cloud [24].

2.2 Related Work

While the problem at hand has not extensively studied, there exist relevant work in the areas of general user security attitudes, usability of security warnings, and traffic analysis for IoT networks.

User security attitudes and models: Dourish et al. [10] found that negative attitudes towards computer security dominate user experience. Gross and Rosson [14] found that users are not well supported by technology in managing sensitive data. Grinter et al. [13] discovered that even technically-savvy users find home network management nontrivial. Wash [25] described how users may discard expert security advice based on their models. Nthala et al. [21] found

that users oftentimes leverage skilled acquaintances to keep their network secure. In IoT-specific work, Zeng et al. [26] uncovered a variety of issues in the attitude of users towards smart home security. Zheng et al. [27] found users lack awareness of IoT-related privacy risks. Emami-Naeini et al. [11] found that users are consistently concerned about IoT privacy/security after device purchase. Overall, these works consistently found that the interface between users and computing devices (particularly IoTs) is problematic, which is one of the motivating arguments for our work.

Usable warnings: Akhawe and Porter-Felt [1] highlighted how the formulation of browser security warnings impacts their effectiveness. Anderson et al. [2] and Bravo-Lillo et al. [6] studied the effect of habituation to users' attention in the context of traditional computing devices. However, the results of this line of research cannot be directly ported to the IoT domain, where significantly different interfaces and modes of interaction may be used.

IoT Traffic analysis: We summarize here a few recent representative works. Bezawada et al. [4] propose a fingerprinting technique to identify device type from network traffic. Miettinen et al. [18] leverage fingerprinting to place devices in different trust classes. The Princeton IoT Inspector project maintains a IoT fingerprinting and browser- based traffic inspection tool [15]. Mirsky et al. [19] describe ensembles of autoencoders as an effective unsupervised anomaly detection technique for the IoT domain. Martin et al. [17] describe a system to identify suspicious traffic and build signatures. While all these works provide an important algorithmic foundation for IoT traffic analysis, they offer limited considerations of the issues arising when attempting to interface with the user.

3 The Case for Involving the User

Cranor, in proposing a framework for integrating humans into security systems, states that a secure system should either keep humans out of the loop, or—if not possible—maximize the chances that the human performs security function successfully [9]. We therefore begin by posing the following question: *can a network intrusion detection system for home IoT devices work fully autonomously?*

We believe the answer is "no". The first reason is Axelsson's insight [3] that most IDS's generate more false positives than true positives. As attacks generally constitute rare events, this is likely to affect even an exceptionally accurate detector. As a case-study, consider the Kitsune detector [19], a state-of-art anomaly-based IDS. An important parameter for a detector is the Bayesian detection rate, i.e., the probability that an alert corresponds to a true attack. The Kitsune authors characterize their algorithm on nine sample attacks, from which the Bayesian detection rate can be computed. Oversimplifying, assume that each attack happens once a day and lasts ten minutes, and that traffic is analyzed in discrete time windows of 1 minute. Under these assumptions, Kitsune's Bayesian detection rate oscillates between 65% and 0.4%, with a median of 43%. While

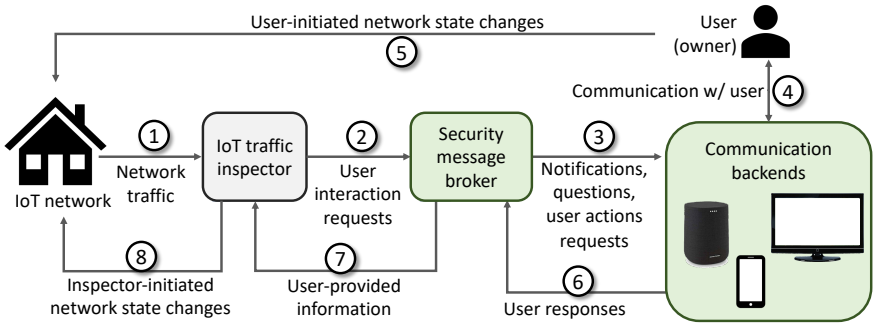

Fig. 1: Overview of proposed architecture

impressive, these results are not sufficient to allow a system to autonomously decide which flows should be blocked, and which ones should be allowed.

Further, a fully automated system would violate what we believe to be an important tenet in the design of human-centered technology: that the user should retain *control* of their devices, and *understanding* of their action. The relationship between users and the security of consumer devices has long been known to be problematic. Over time, studies on user perception of computer security have found attitudes that express frustration [10], concern [11], difficulty [13], and the necessity to rely on expert friends and relations [21]. In the IoT world, these difficulties are worsened by the deeply invasive nature of security breaches (e.g., access of video/audio from within the home). A device making autonomous, unexplainable, and oftentimes incorrect security decisions carries the risk of further aggravating the situation, by causing unexplained malfunctions (e.g., shutting a Smart TV off the network due to false positives).

Finally, we briefly remark that, even ignoring the two considerations above, securing IoTs occasionally requires manual operations to perform configuration changes, firmware updates [23], reboots, etc.

4 Proposed Architecture and Challenges

Based on the discussion in Section 3, we believe that systems that do not involve the user are not viable. In order to productively discuss the challenges of a system that does involve the user, we now propose a general architecture.

4.1 Architecture

We envision an IoT residential network incorporating a monitoring components to identify misbehaving and/or insecure devices, and communicate/interact with users to resolve security-critical situations. The expected system architecture is summarized in Figure 1. An IoT inspector component analyzes local traffic ①. The analysis may result in various *user interactions*. These include:

1. **Notifications**, in which a security finding is communicated to the IoT user.
2. **Questions**, through which the inspector gathers additional context (e.g., asking the user to describe recent activity to disambiguate findings).
3. **Action requests**, through which the inspector requires manual-only operations (e.g., manual factory reset of a compromised router).

All such interactions are directed towards the user through a message broker ②. This component is necessary as, in a home network, the user can generally be reached using mobile apps, or the IoT devices themselves. It is therefore necessary to determine the best device, among those available, to perform the requested user interaction. The broker performs this function ③. As a result of the communication issued through devices ④, the user may perform actions on the network directly ⑤ (e.g., rebooting a router), or respond to questions ⑥. User responses are received and parsed by the broker, and their content is provided to the inspector ⑦. The inspector may also attempt to automatically resolve issues based on user feedback ⑧.

4.2 Scenarios

In order to highlight the challenges of bringing the user in the loop of a IoT security system, we briefly describe two scenarios. These take place in a hypothetical households which includes a homeowner and a number of smart devices.

Scenario #1: privacy violation. The traffic inspector detects unusual bursts of activity from an adjustable home monitoring camera in the living room. Given that the smart speaker and TV are currently off, the message broker issues a notification towards the user smartphone. Upon confirmation that the user is not operating the camera, the inspector informs the user that the camera may have been hacked. The user inspects the camera and realizes that it appears to be reorienting itself without anyone in the household controlling it. The user suspects this event is related to an abusive former partner and disables the camera.

Scenario #2: malware break-in. The traffic inspector generates a notification as the smart fridge is suddenly emitting denial-of-service traffic. It detects that the device is running old firmware with an hardcoded-credential vulnerability. As the smart speaker is currently being used, the broker directs the speaker to enunciate a spoken warning. The warning directs the user to a more extensive smartphone notification which includes instructions on (i) how to reset the appliance's firmware, and (ii) how to perform a step-by-step firmware update.

None of these scenarios is futuristic or far-fetched. All elements, including commandeering of internet-connected cameras [22], vulnerable home appliances [7], the use of IoT devices for domestic abuse [5], hardcoded-credential vulnerabilities [16] and IoT-based DDoS [12] have been observed in the wild. Likewise, research in the IoT space has looked at anomaly detection [19], vulnerability identification [18], and semi-automated firmware updates [23].

4.3 Challenges and Opportunities

The scenarios above highlight a number of open questions which concern the specifics of how warning and interactions are phrased and communicated, including: *(i) Do warnings communicated using different modes (e.g., audio vs text) receive the same attention? (ii) How can we choose the best mode of interaction depending on context? (iii) How can security notifications be formulated to ensure the user understands the situation without receiving undue stress? (iv) How should warnings be managed in a multi-user household? (v) Do warnings and interactions need to be adjusted based on demographics (e.g., age group)?* These questions suggest that, for a network monitoring system, productively cooperating with the user is nontrivial. They also suggest some general directions for future research. First, designers of IoT network monitoring systems should focus on the explainability of their findings (intended as the ability to correlate low-level network behavior, e.g., an increase in short-lived flows, to high-level concepts, e.g., a DDoS). Second, the heterogeneity of IoT devices creates an opportunity to explore user interactions via a variety of modes and interfaces (audio/visual signals, text notifications, haptic, etc.).

5 A Testbed for Evaluating IoT Security Warnings

As an initial exploration in user-interactive IoT monitoring systems, we designed and instantiated a testbed for user studies involving IoT-related security warnings. The design aims at easing common practical issues that affect IoT security and usability research. Not every institution can afford to deploy a dedicated IoT lab. In some case, dedicated space for user studies may not be available, or it may be shared between multiple projects (not all IoT-related). Acquiring the devices may itself be a problem due to the high price of some of them.

Our testbed design implements the components shaded in green in Figure 1. The software powering the testbed consists of a simple message broker and notification backends for three popular types of IoT devices (we make the software publicly and freely available[1]). We specifically aimed at supporting devices (Android phone; Sonos smart speaker; Raspberry Pi-based Smart TV) which are cheap to acquire and can be deployed and dismantled in a matter of minutes. We hope releasing our testbed software will foster its deployment in research labs, HCI classes, REUs, and other scenarios.

The focus of our testbed is on studying how users receive and understand interaction requests from an IoT monitoring system, and how the phrasing and the mode (audio vs text) of these notifications affect comprehension. The intended mode of use is to have subjects sit through simulated security incidents, where an experimenter generates notifications and questions towards the user. The testbed does not include or propose novel intrusion detection/monitoring algorithms, although such components could easily be interfaced with the broker.

[1] `https://www.dropbox.com/sh/etkdrmddg2ilkqr/AAC_rctdyOShzMlJu-kCFJoGa`

6 A Preliminary User Study

We conclude this paper by describing a qualitative, IRB-approved user study which we conducted to test-drive the testbed and receive feedback. Recruiting was conducted via on-campus advertising and word-of-mouth. Overall, we recruited 10 participants (all college students from our metro area).

6.1 Experiment Design

Each subject participated in an individual network security exercise. The exercise simulates two different scenarios, both involving an intrusion in the home network to which the testbed devices are connected. The experimenter generates interactions with the participant in order to resolve the issue. Both scenarios are representative of issues where an attacker attempts to masquerade communication to/from an infected device as video streaming from phone or TV.

Prior to the experiment, each subject was informed that during the experiment they would receive notifications about a simulated IT security incident, and they would be requested to pass information to the system. At the end of the experiment, each subject was asked to complete a brief questionnaire.

6.2 Results

Due to the limited sample size and diversity, we do not consider the study sufficient for drawing general conclusions, and we omit a quantitative discussion of the results. However, the study uncovered several interesting user attitudes that highlight some research challenges. We report a few relevant findings.

The most relevant feedback was provided in response to the following question: *How likely would you be to use a system like the one you interacted with in your own home?*. The question required an answer between 1 ("not at all") and 5 ("a lot"), and received the following answers: 3,3,4,2,3,4,4,5,5,3. It further gave each subject the opportunity to provide an unstructured explanation of their answer. Participants who answered in the 4-5 range tended to underscore the importance of keeping the network secure. They also expressed caveats (Subject #6: "How likely I would be to use it is proportional to how accurate it is"). Participants in the range 1-3 expressed, among other concerns, issues with frequent interruptions (Subject #4: "If I was frequently alerted to IT security issues, I would simply stop using the internet."). This is consistent with other user studies that found that convenience tends to trump privacy concerns [27]. Others expressed the concern that the system itself may be prone to security issues (Subject #2: "the notification method itself is prone to security issues").

When asked *What do you think can be improved?*, Subject #10 stated "something that "looks nice" lends it a more professional apperence [sic], and therefore an implicit level of trust on behalf of the user.". Similarly, Subject #6 asked for "An authentication process.". Furthermore, Subject #1 asked "May be [sic] just simple language to define the issues so that even a novice user could understand without having a better network knowledge", and Subject #4 stated "Have the notifications be more educational".

6.3 Discussion

The study found confirmation of the usual hurdle to adoption of a security system: the need to minimize disruption. The issues of security and trust (Subjects 2, 6, 10) are important, as an unprotected notification system may be abused by an attacker. They also suggest that it may be necessary to visually distinguish authenticated notifications. The complaints concerning language (Subjects 1, 4) highlight the difficulty of appropriately phrasing interaction requests. Overall, user feedback suggests that there is interest for technologies, that can help users to identify and understand security issues in their home networks.

7 Conclusion

In this paper, we advocated the need for research on IoT security techniques which are aware of, and involve the user. In particular, we claim that, since most IoT device owners are not security experts, it is necessary to devise novel forms of communication of security findings which are accessible to this population. In order to foster further work in the area, we contributed a general architecture for IoT network monitoring tools, which keeps the user in the loop. Furthermore, we presented the design and implementation of a testbed for research on user-facing network security notifications and feedback requests.

References

1. Akhawe, D., Felt, A.P.: Alice in Warningland: A Large-Scale Field Study of Browser Security Warning Effectiveness. In: USENIX Security Symposium. p. 17 (2013)
2. Anderson, B.B., Kirwan, C.B., Jenkins, J.L., Eargle, D., Howard, S., Vance, A.: How polymorphic warnings reduce habituation in the brain: Insights from an fmri study. In: CHI (2015)
3. Axelsson, S.: The base-rate fallacy and its implications for the difficulty of intrusion detection. In: Proceedings of the 6th ACM conference on Computer and communications security - CCS '99. pp. 1–7. ACM Press (1999)
4. Bezawada, B., Bachani, M., Peterson, J., Shirazi, H., Ray, I., Ray, I.: Behavioral Fingerprinting of IoT Devices. In: ASHES. New York, NY, USA (2018)
5. Bowles, N.: Thermostats, locks and lights: Digital tools of domestic abuse (2018), https://www.nytimes.com/2018/06/23/technology/smart-home-devices-domestic-abuse.html
6. Bravo-Lillo, C., Cranor, L., Komanduri, S., Schechter, S., Sleeper, M.: Harder to ignore? In: SOUPS (2014)
7. Chirgwin, R.: Dishwasher has directory traversal bug (Mar 2017), https://www.theregister.com/2017/03/26/miele_joins_internetofst_hall_of_shame/
8. Cimpanu, C.: Hacker leaks passwords for more than 500,000 servers, routers, and IoT devices (Jan 2020), https://www.zdnet.com/article/hacker-leaks-passwords-for-more-than-500000-servers-routers-and-iot-devices/
9. Cranor, L.F.: A Framework for Reasoning About the Human in the Loop. In: UPSEC (2008)

10. Dourish, P., Grinter, E., Delgado de la Flor, J., Joseph, M.: Security in the wild: User strategies for managing security as an everyday, practical problem. Personal Ubiquitous Comput. **8**(6), 391–401 (Nov 2004)
11. Emami-Naeini, P., Dixon, H., Agarwal, Y., Cranor, L.F.: Exploring how privacy and security factor into iot device purchase behavior. In: CHI (2019)
12. Goodin, D.: Record-breaking ddos reportedly delivered by ¿145k hacked cameras (2016), https://arstechnica.com/information-technology/2016/09/botnet-of-145k-cameras-reportedly-deliver-internets-biggest-ddos-ever
13. Grinter, R.E., Edwards, W.K., Chetty, M., Poole, E.S., Sung, J.Y., Yang, J., Crabtree, A., Tolmie, P., Rodden, T., Greenhalgh, C., Benford, S.: The ins and outs of home networking: The case for useful and usable domestic networking. ACM Trans. Comput.-Hum. Interact. **16**(2), 8:1–8:28 (Jun 2009)
14. Gross, J.B., Rosson, M.B.: Looking for trouble: Understanding end-user security management. In: CHIMIT (2007)
15. Huang, D.Y., Acar, G., Apthorpe, N., Li, F., Narayanan, A., Feamster, N.: Princeton IoT Inspector (May 2019), https://iot-inspector.princeton.edu
16. Lewellen, T.: CERT/CC Vulnerability Note VU#800094 (Sep 2013), https://www.kb.cert.org
17. Martin, V., Cao, Q., Benson, T.: Fending off IoT-hunting attacks at home networks. In: CAN (2017)
18. Miettinen, M., Marchal, S., Hafeez, I., Asokan, N., Sadeghi, A., Tarkoma, S.: IoT SENTINEL: Automated Device-Type Identification for Security Enforcement in IoT. In: ICDCS (2017)
19. Mirsky, Y., Doitshman, T., Elovici, Y., Shabtai, A.: Kitsune: An Ensemble of Autoencoders for Online Network Intrusion Detection. In: NDSS (2018)
20. Nichols, S.: Don't panic, but your baby monitor can be hacked into a spycam (Jun 2018), https://www.theregister.co.uk/2018/06/22/baby_monitor_hacked/
21. Nthala, N., Flechais, I.: Informal support networks: an investigation into home data security practices. In: SOUPS (2018)
22. Pascu, L.: Multiple critical security flaws found in nearly 400 IP cameras - Bitdefender BOX Blog (Jun 2018), https://www.bitdefender.com/box/blog/ip-cameras-vulnerabilities/multiple-critical-security-flaws-found-nearly-400-ip-cameras/
23. Simpson, A.K., Roesner, F., Kohno, T.: Securing vulnerable home IoT devices with an in-hub security manager. In: PerCom Workshop (2017)
24. Taylor, C.R., Shue, C.A., Najd, M.E.: Whole home proxies: Bringing enterprise-grade security to residential networks. In: IEEE ICC (2016)
25. Wash, R.: Folk models of home computer security. In: SOUPS (2010)
26. Zeng, E., Mare, S., Roesner, F.: End user security and privacy concerns with smart homes. In: SOUPS (2017)
27. Zheng, S., Apthorpe, N., Chetty, M., Feamster, N.: User perceptions of smart home iot privacy. Proc. ACM Hum.-Comput. Interact. **2**(CSCW), 200:1–200:20 (Nov 2018)