Network-Level Security for the Internet of Things: Opportunities and Challenges

Hassan Habibi Gharakheili, Arunan Sivanathan, Ayyoob Hamza, and Vijay Sivaraman, University of New South Wales

Smart environments with many Internet of Things (IoT) devices are at significant risk of cyberattacks, putting private data and personal safety in danger. While IoT device manufacturers are putting more safeguards in their products, they need to be augmented with network-level methods to detect and block anomalous behavior. Our approach provides a strong layer of runtime defense at the network layer applicable to large and heterogeneous IoT environments.

Several research groups[1,2] have identified and reported various forms of vulnerabilities in commercial Internet of Things (IoT) devices. We experimentally evaluated the privacy and security risks of tens of consumer IoT devices[4] and demonstrated real-life threats to typical users posed by hackers: they can snoop on activities inside a building by obtaining unencrypted data from motion detectors and security cameras or disable the smoke alarms by bombarding them with a large number of requests. Furthermore, most commercial IoT devices can be exploited to reflect and amplify attacks on Internet-based servers, while certain models of connected printers allow hackers to obtain recently scanned documents or print threatening messages remotely. A systematic technique[3] for evaluating the security posture of IoT devices can be divided into four categories based on 1) the confidentiality of data they communicate, 2) the integrity of connections they make, 3) the access control of devices, and 4) their capacity in

Digital Object Identifier 10.1109/MC.2019.2917972
Date of publication: 30 July 2019

FROM THE EDITOR

Standard security parlance often adopts a trust-and-verify model for working with systems. However, many emerging Internet of Things ecosystems will often encompass untrusted devices that pose potential security risks. Rather than simply assuming that a device behaves as expected, it then becomes incumbent on the supporting infrastructure to defend the complete system against attack. In this article, the authors consider the specific role of the network in securing a system to provide layers of defense without inherently trusting the device. Ranging from a detailed analysis of the current behavior of devices to analysis and enforcement mechanisms for handling network flows, this article provides a solid end-to-end perspective on a secure IoT network. — *Trevor Pering*

reflecting unwanted traffic to other Internet-connected services. Table 1 shows sample outcome ratings obtained for three devices. The smoke alarm properly protected confidentiality (green boxes), the camera had poor access control protection (red boxes), firmware updates were manual rather than automatic for the camera and motion sensor (yellow boxes), and some attributes could not be tested or assessed (gray unfilled boxes).

BASELINE BEHAVIOR OF IOT DEVICES IN THE NETWORK

Network operators today lack real-time visibility into connected devices: more than 40% of today's endpoints are unknown and unmanaged by their organizations, leading to significant infrastructure blind spots, unauthorized access, and data leaks.[5] Given that IoT devices exhibit limited traffic patterns, we believe it is possible to identify and profile their network behavior. The network traffic of tens of real IoT devices over a period of several months[6] can be represented pictorially using Sankey diagrams, as shown in Figure 1, for two representative IoT devices: a LIFX smart bulb and an Amazon Echo. Using these representations, we developed a framework[7] that is able to classify IoT devices based on the statistical attributes of their network traffic, such as activity cycles and volume patterns, server-side port numbers, Domain Name System (DNS)/Network Time Protocol (NTP) signaling profiles, and cipher suites exchanged in Secure Sockets Layer connections. This enables network operators to have runtime visibility and the baseline behavior of operational IoT devices in their network.

TABLE 1. The security posture rating for selected IoT devices.

Devices	Confidentiality										Integrity and Authentication					Access Control							Reflection			
	Device to Server			Device to Application			Application to Server			All																
	Plaintext	Protocol	Entropy	Plaintext	Protocol	Entropy	Plaintext	Protocol	Entropy	Privacy	Replay Attack	DNSSEC	DNS Spoofing	Fake Server	Firmware	Open Ports (TCP)	Open Ports (UDP)	Vulnerable Ports	Weak Passwords	ICMP DoS	UDP DoS	Number of TCP Connections	ICMP Reflection	SSDP Reflection	SNMP Reflection	SNMP Public Community String
Camera	A		A	C	C	A	A	A	A	C	A	C	C	A	B	C	C	C	C	C	B	C	C	A	A	A
Motion Sensor	A	A	A	C	C	C	A	A	A	C	A				B	C	C	A	A	C	B	C	C	C	A	A
Smoke Alarm	A		A	A	A	A	A	A	A	A	A	C	C	A	A	B	C	A	A			A	C	A	A	A

DNSSEC: Domain Name Security Extensions; TCP: Transmission Control Protocol; UDP: User Datagram Protocol; ICMP: Internet Control Message Protocol; DoS: denial of service; SSDP: Simple Service Discovery Protocol; SNMP: Simple Network Management Protocol.

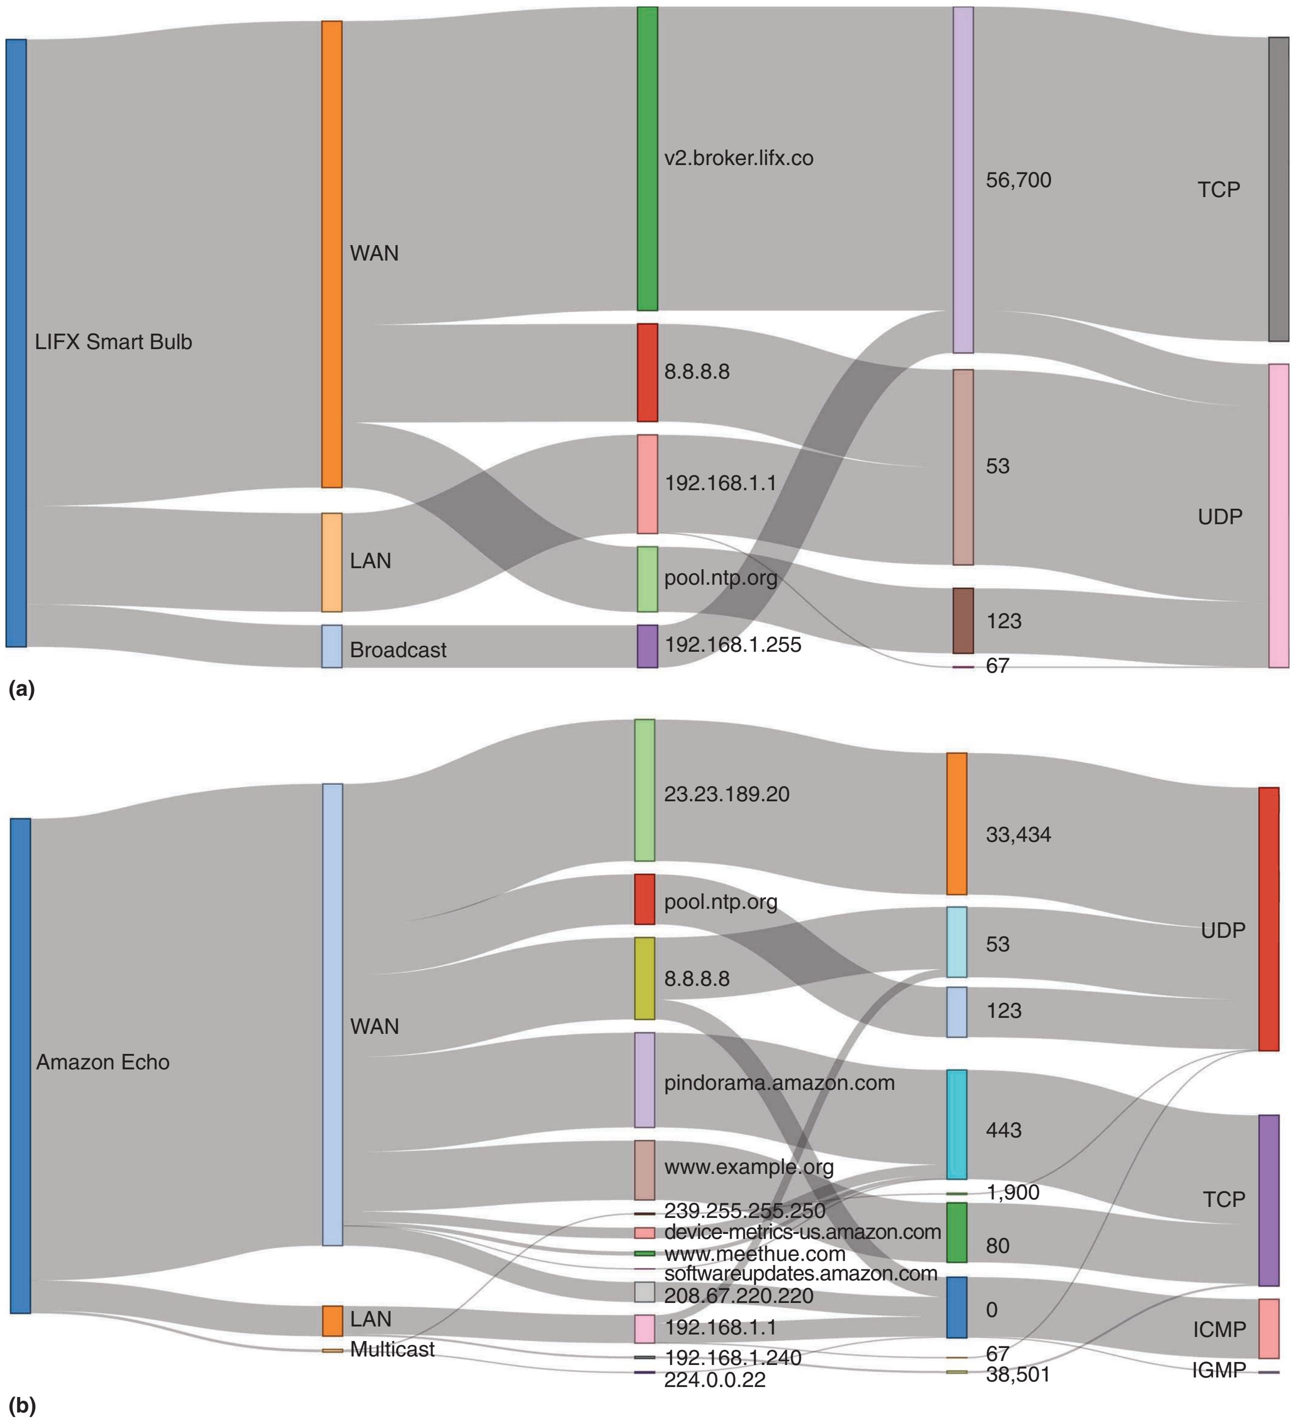

FIGURE 1. A Sankey diagram of daily network activity for two representative IoT devices: (a) a LIFX smart bulb and (b) an Amazon Echo. LAN: local area network; WAN: wide area network; IGMP: Internet Group Management Protocol.

FORMAL BEHAVIORAL MODELS AND RUNTIME LOCKDOWN

A formal grammar for IoT network behavior in the form of access control lists (ACLs) was recently standardized by the Internet Engineering Task Force in the form of RFC Standard 8520, *Manufacturer Usage Description* (MUD). MUD allows the manufacturer to specify what layer 3 servers and layer 4 ports may or may not be used by the device. For example, an Internet Protocol camera may need to use DNS and dynamic host configuration protocol on the local network, contact NTP servers, and communicate over HTTPS with a set of cloud-based controllers on the Internet but nothing else. Knowing each device's requirements allows network operators to impose a tight set of ACL restrictions for each IoT device in operation to reduce the potential attack surface. We developed a tool that can automatically generate the MUD profile for an IoT device from its (presumably clean) traffic trace and further translate the MUD profile into flow rules that can be enforced at runtime into the network substrate using the software-defined networking (SDN) paradigm, as shown in Figure 2. This lockdown of the IoT device ensures that conformant traffic can be let through without inspection, while unexpected packets can be either dropped or passed for inspection to an intrusion detection system to detect security breach attempts.[8]

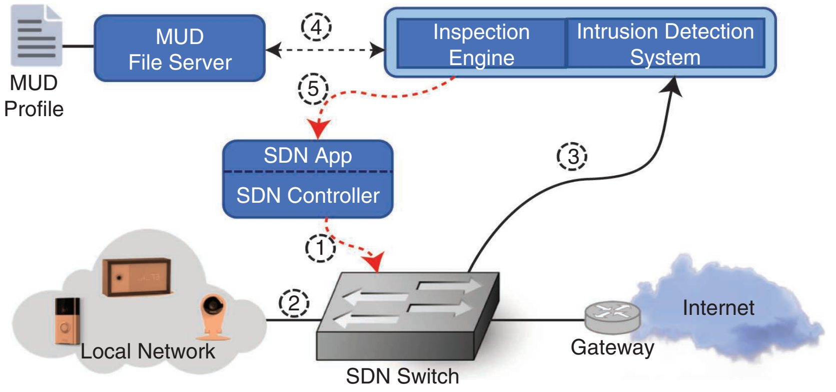

FIGURE 2. An automatic translation of MUD policies and enforcement of network flow rules.

REAL-TIME ANOMALY DETECTION

The enforcement of MUD-policy-based lockdown using SDN can significantly reduce the attack surface on the IoT device, but certain volumetric attacks can still be launched. We therefore developed methods[9] for detecting volumetric attacks that are not prevented by the MUD profile because ACLs simply allow or deny traffic without provision to limit rates. Fending off such attacks requires more sophisticated machinery (shown in Figure 3). This machinery monitors the level of activity associated with each policy rule compared to the baseline (deduced in the "Baseline Behavior of IoT Devices in the Network" section) to detect anomalies. Scalable models for detecting anomalous patterns can be trained by intelligently combining coarse-grained (device-level) and fine-grained (flow-level) flow telemetry, and this continues to be the subject of our ongoing research.

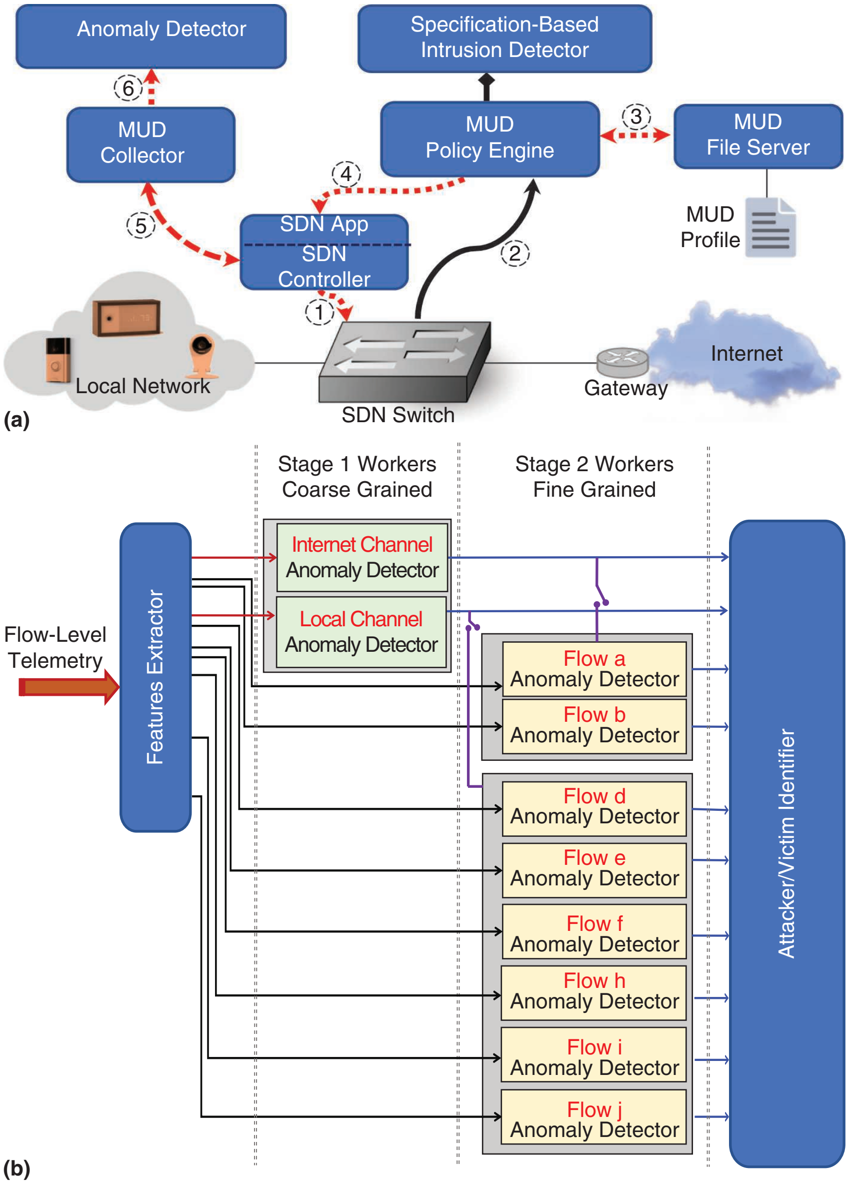

FIGURE 3. Detecting anomalies by real-time-monitoring, MUD-compliant network flow rules: (a) an SDN-based system architecture and (b) a device-specific anomaly detection.

IoT security research is still in its early days, but it is becoming clear that device-level security will by itself not suffice and network-level protection measures will be needed. In this article we outlined our approach, which combines baseline network behavior with SDN lockdown and anomaly detection to identify and eventually mitigate security attacks on IoT-rich environments. In the big picture of IoT system security, these steps form the foundation for smarter and more robust networks that will enable more complex deployments. Without such systems, smart environments will remain at constant risk to unauthorized intrusions. Thus far, our approach, which has been very much device-centric, needs to be augmented by network-wide monitoring to capture interrelationships among IoT devices, users and servers, and controllers in smart environments.

REFERENCES

1. N. Dhanjani, *Abusing the Internet of Things: Blackouts, Freakouts, and Stakeouts*. Sebastopol, CA: O'Reilly Media, 2015.
2. E. Fernandes, J. Jung, and A. Prakash, "Security analysis of emerging smart home applications," in *Proc. 2016 IEEE Symp. Security and Privacy (SP)*, pp. 636–654.
3. F. Loi, A. Sivanathan, H. H. Gharakheili, A. Radford, and V. Sivaraman, "Systematically evaluating security and privacy for consumer IoT devices," in *Proc. 2017 Workshop on Internet of Things Security and Privacy*, pp. 1–6.
4. A. Sivanathan, F. Loi, H. H. Gharakheili, and V. Sivaraman, "Experimental evaluation of cybersecurity threats to the smart-home," in *Proc. 2017 IEEE Int. Conf. Advanced Networks and Telecommunications Systems (ANTS)*, pp. 1–6.
5. Cisco, "Cisco 2017 midyear cybersecurity report," 2017. [Online]. Available: https://www.cisco.com/c/dam/global/es_mx/solutions/security/pdf/cisco-2017-midyear-cyber security-report.pdf
6. UNSW Sydney IoT Security. Accessed on: June 13, 2019. [Online]. Available: https://iotanalytics.unsw.edu.au/index
7. A. Sivanathan et al., "Classifying IoT devices in smart environments using network traffic characteristics," *IEEE Trans. Mobile Comput.* doi: 10.1109/TMC.2018.2866249.
8. A. Hamza, H. H. Gharakheili, and V. Sivaraman, "Combining MUD policies with SDN for IoT intrusion detection," in *Proc. 2018 Workshop on IoT Security and Privacy*, pp. 1–7.
9. A. Hamza, H. H. Gharakheili, T. A. Benson, and V. Sivaraman, "Detecting volumetric attacks on IoT devices via SDN-based monitoring of MUD activity," in *Proc. 2019 ACM Symp. SDN Research*, pp. 36–48.

HASSAN HABIBI GHARAKHEILI is with the University of New South Wales, Sydney, Australia. Contact him at h.habibi@unsw.edu.au.

ARUNAN SIVANATHAN is with the University of New South Wales, Sydney, Australia. Contact him at a.sivanathan@unsw.edu.au.

AYYOOB HAMZA is with the University of New South Wales, Sydney, Australia. Contact him at m.ahamedhamza@unsw.edu.au.

VIJAY SIVARAMAN is with the University of New South Wales, Sydney, Australia. Contact him at vijay@unsw.edu.au.

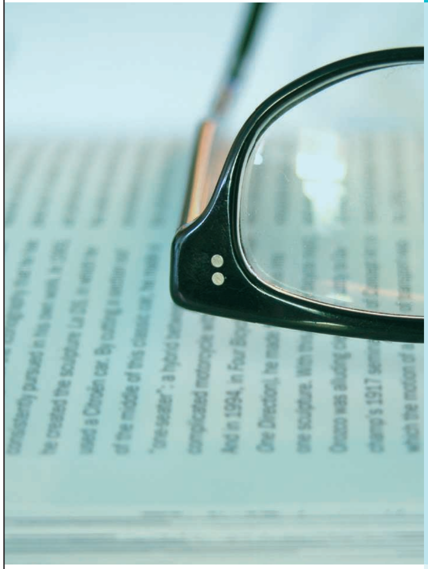

Digital Object Identifier 10.1109/MC.2019.2926564