Sounding the Bell for Improving Internet (of Things) Security

Theophilus Benson
Brown University
theophilus_benson@brown.edu

Balakrishnan Chandrasekaran
Technische Universität Berlin
balac@inet.tu-berlin.de

ABSTRACT

The fragility of the Internet of Things (IoT) ecosystem poses serious threats to Internet security, and the proliferation of IoT devices only exacerbates this situation by providing vulnerable end-points to be exploited and used as attack sources. While industry and academia are working hard on designing innovative solutions to detect, mitigate and thwart massive botnet-based DDoS attacks, the space of solutions appears disjoint and fragmented. The lack of cooperation between the IoT device manufacturers, network operators, content providers, end users, and other players precipitates in point solutions which offer at best a veneer of security. In this paper we alert the community to the security challenges posed by the fragile IoT ecosystem, discuss the space of solutions, and present the need for a distributed, concerted effort, e.g., among end users, ISPs, and CDNs, to improve Internet security. We do not claim to solve the problem, but offer design guidelines and discuss the key implementation challenges to inform the debates on IoT security.

CCS CONCEPTS

• **Networks** → **Denial-of-service attacks**; *Network privacy and anonymity*; • **Security and privacy** → *Security protocols*;

KEYWORDS

Internet of Things (IoT), Botnet, DDoS attacks

1 INTRODUCTION

The Internet was designed with an implicit notion of trust, and as such has been subjected to a broad spectrum of security threats and attacks during the last decade. While there exists an extensive body of prior work on improving the resilience of the Internet to various forms of security threats, the battle is far from over. Attacks of more than 300 Gbps, for instance, have become more common [2, 3] than before. The unprecedented growth of Internet-of-Things (IoT) devices [17, 19], and the emergence of a connected *Internet of Everything* coupled with the fragility of the IoT ecosystem [7, 16, 21, 29, 34, 42] provides a ripe platform for inimical parties to exploit the weaknesses, and launch massive DDoS attacks. In this regard, the recent Mirai botnet attacks [27, 35] are only a harbinger of more widespread and crippling attacks in the future.

IoTS&P'17, November 3, 2017, Dallas, TX, USA

ACM ISBN 978-1-4503-5396-0/17/11…$15.00
https://doi.org/10.1145/3139937.3139946

While the security community is all too familiar with issues of bugs and unpatched software, insecure defaults, weak authentication mechanisms, and poorly-configured systems, the emergence of the IoT ecosystem exacerbates these issues. The sheer volume of IoT devices in the market—expected to reach 8.4 billion by the end of 2017, outnumbering the world's population by almost 1 billion [19]—not only amplifies the well-known security issues, but renders current solution strategies ineffective in mitigating and controlling botnet-based DDoS attacks. Although several point solutions have been proposed [1, 30, 36], we argue that these are precisely just that; these point solutions offer, if any, only a veneer of security, and we are in a dire need of a more distributed and concerted approach to securing the Internet.

The recent Mirai botnet attacks [27, 35] generated attack traffic of unprecedented volumes, even making, in one scenario, the DDoS mitigation service offered by a cloud provider economically unviable [31].[1] The attack garnered widespread media attention [24, 28, 32, 45], and led to a global realization of the need for regulations concerning the design, implementation, and deployment of IoT devices [10, 44]. Using this attack as a case study, we highlight the key challenges in securing the Internet against such large-scale botnet-based DDoS attacks.

In this paper we take the position that detecting, mitigating and avoiding large-scale botnet attacks relies on a highly distributed, but concerted effort of all interested stakeholders, e.g., the device manufacturers, consumers or end users, network operators, content delivery networks (CDNs), and cloud service providers. We share a vision of a framework that exploits the Internet's hierarchical topology to realize a practical, scalable security solution. In sharing this vision, we outline key design guidelines, and discuss the inherent tradeoffs that we must balance in realizing an implementation. We summarize our contributions as follows.

* We highlight the fragility of the IoT ecosystem, and discuss how the Mirai botnet exploited this ecosystem to stage a large-scale DDoS attack.

* We share a few key insights into the Mirai botnet attacks that reveal the inherent challenges in detecting, mitigating, and avoiding such large-scale botnet attacks in the future.

* We outline the design of a security architecture that exploits the hierarchical nature of the Internet's architecture to secure the Internet. Our solution relies on the cooperation between the Internet's stakeholders, e.g., consumers, network operators, and IoT device manufacturers.

* We discuss the fundamental tradeoffs that we must balance in realizing a practical, scalable security solution.

[1] We note, however, that the DDoS mitigation service was offered as a pro bono service by the cloud provider.

Organization In §2 we discuss the fragility of the IoT ecosystem and share key insights into the Mirai botnet attacks. We discuss, in §3 the key challenges inherent in detecting and mitigating large-scale botnet attacks. §4 presents an outline of a new security architecture, and explains the tradeoffs we must balance during the implementation. We briefly review the related work in §5, and present our conclusions in §6.

2 BACKGROUND

The number of IoT devices is expected to reach 8.4 billion by the end of 2017 [19]. But, with IoT device manufacturers trying to outdo one another in terms of device features, security has been relegated to the back seat. Consequently, the staggering growth of the IoT ecosystem should primarily be a huge *cause for concern* to security researchers.

2.1 A fragile ecosystem

While the problems of unpatched software, bugs and zero-day vulnerabilities, and misconfigurations are generally well-known, the sheer volume of the IoT devices amplifies the impact of such well-known issues by a huge factor. To examine how and why the IoT ecosystem exacerbates security issues, we highlight a few new security issues that the IoT devices pose.

Insecure, unpatchable IoT devices often lack a streamlined procedure for patching vulnerabilities, and some are also virtually unupdatable. The lack of update mechanisms combined with a widespread deployment creates an ecosystem that is ripe for exploitation by inimical parties [26, 29, 33].

Lack of a user interface IoT devices typically lack a user or management interface for updating or reviewing the device configuration. Without a user interface the IoT devices cannot signal the end user of the availability of a critical patch or security update, or alert the user of suspicious activity.

Misconfigurations and insecure defaults IoT devices are often bought, deployed at home, and operated as if they are configured *a priori* with the best security defaults. Configuration defaults, however, are worst in practice: simple "hard coded" credentials, insecure administration panels, exposed ports and services are rampant among deployed devices [7, 12, 13, 18, 29].

Misplaced trust Although the IoT devices are gaining more computing power, the support for monitoring or policing these devices is not even an afterthought; they are hardly even considered as part of the design guidelines. Installing antivirus software, for instance, is not possible, even though the compromise of many types of IoT devices, e.g., Nest Thermostat, could potentially endanger the end users.

Unregulated ecosystem Despite the lack of options for updating or patching most IoT devices, and poor configuration defaults rendering any notion of security obsolete, debates on devising regulations governing the need for securing the IoT ecosystem are far and few between. The few proposals for regulating IoT security that are on the horizon [10, 44] still leave several key questions unanswered; the ability to patch or update does not automatically translate to up-to-date devices, since users have little or no incentive to update devices.

In the remainder of this section, we briefly review the Mirai botnet attacks and summarize the challenges in designing a practical and scalable solution to secure the IoT ecosystem.

2.2 Exploiting the fragile ecosystem

The recent, unprecedented DDoS attacks from the Mirai botnet [27, 35] present an interesting case study to obtain insights into the security threats posed by the IoT ecosystem. Rather than describe the botnet attacks, we selectively present the key observations that highlight how the botnet exploited the fragility of the IoT ecosystem and reveal the inherent challenges in securing the Internet.

Widespread attack sources The Mirai botnet is one of the largest botnets ever encountered by the Internet, boasting a network of some 600, 000 devices at some point in its evolution (refer Figure 3 in [5]). The IoT devices in the Mirai botnet spanned a wide geographic footprint and were from several different networks.

Low-barrier to hacking An appalling observation on the Mirai botnet is that the botnet gained control of hundreds of thousands of devices using just 62 credentials [13], some of which are factory *defaults* that are left unchanged by users.

Inconspicuous growth The growth rate of the Mirai botnet during its initial or *bootstrap* phase was much smaller than that of others. Regardless of the conjectures on the slow growth [5], we note that the slow bootstrap phase did not deter the efficacy of botnet's attacks significantly. This observation perhaps underscores the lack of a "need" for fast growth given that the Mirai botnet eventually did capture a large number of devices.

Low per-device traffic Though the botnet generated an unprecedented volume of traffic, e.g., over 1 Tbps of attack traffic against the French hosting provider OVH [27], the per-device attack traffic from most devices was extremely low [6].

Unprecedented attack volumes The approximately 600 Gbps of traffic in case of the attack against "Krebs on Security" blog, and the 1.1 Tbps of traffic in the attack against the French hosting provider OVH were not generated when the botnet reached its peak size. In either case the number of devices contributing to the attack were far lower compared to the 600, 000 devices that the Mirai botnet controlled at its peak growth. The attacks we have witnessed are yet far from the what perhaps is potentially feasible with such botnets.

Shape-shifting attacks A critical threat posed by today's botnets is that the large-scale attacks they facilitate do not conform to any particular form or behavior. The recent *pulse wave* attacks, for instance, have no visible ramp-up period: The attackers were able to generate traffic of 300 Gbps or more within a matter of seconds [20, 43]. The attack traffic also exhibited an *on-off* cycle, repeated over a long term. Security researchers claim that rather than wasting time during of the *off*-periods, the attackers switched the targets (victims) on the fly. Such novel attack methodologies necessitate the need for a fast, scalable, and rigorous methods for detecting, and mitigating future large-scale DDoS attacks.

3 DESIGN CHALLENGES

To offer guidelines into designing secure, scalable solutions, we identify the key challenges in instantiating a practical IoT security framework.

Needle in a stack of needles

With widespread attack sources, often spanning several networks and several countries, today's DDoS attacks (e.g., Mirai botnet attacks [5]) require each IoT device only to generate a trickle of the aggregate attack traffic. This low volume of attack traffic from each IoT device makes it highly improbable to detect attacks from close to end users, in home networks where the IoT devices are deployed. Moving away from home networks to a vantage point such as a CDN's infrastructure still does not help much; it is practically infeasible to differentiate "malicious" traffic from the "benign" traffic of a CDN, not at least until the traffic volume crosses a substantial threshold, diminishing the possibility of designing an early warning and DDoS control system at a CDN.

Anomaly detection typically becomes impractical when the anomalies in question are low-probability (or extremely rare) events [39]. Aggregating traffic from multiple sources (or networks) to increase the likelihood of capturing these low-probability events, in contrast, introduces scalability issues; the aggregation points quickly transform into "choke points". For attacks targeting the service-level agreements (SLAs) of a CDN, it suffices to saturate such choke points, making the detection and control of the attack traffic economically unviable for the CDN. The inefficiency of off-the-shelf "learning" algorithms to detect and thwart botnet attacks is a major impediment to designing sophisticated, automated security solutions.

No one vantage point suffices

Each vantage point offers a different perspective into the network traffic. To be practical, effective, and scalable a solution architecture must combine the unique insights gathered from the different vantage points.

Monitoring from edge networks, particularly from within home networks offers an unparalleled perspective: the ability to observe each and every IoT device in close proximity, and potentially "learn" their behavior. A control application running within a home router, for instance, can easily monitor and secure the IoT devices in the home network [36]. While such control applications can continuously monitor the domains or hosts that an IoT device interacts with to check for anomalous behavior, the absence of an established *baseline* behavior for each IoT device limits its applicability. Besides, modifying WiFi router firmware is typically a non-trivial process; mistakes can "brick" the hardware or render the network insecure or inoperable.

Anomalous behavior of one or more devices in any particular home network does not imply a DDoS attack. But such (anomalous) behavior exhibited by devices from a substantial number of homes should raise a red flag. This observation highlights the need for cooperation, for instance, between the control applications in home networks and middleboxes operated by the network provider. Establishing the interfaces for such cooperation is, however, non-trivial. Besides the data exchanged over such interfaces could have privacy implications for the involved parties.

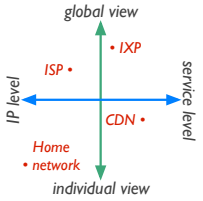

Figure 1: *Four quadrants of the solution space*

Incentives for cooperation

The challenges discussed so far clearly highlight the need for cooperation between stakeholders—device manufacturers, end users, ISPs, IXPs, and CDNs—to instantiate a robust, scalable security framework. Despite the recent crippling DDoS attacks, there is no strong reason to believe that such cooperation between stakeholders will be forthcoming.[2] We need to explore the use of privacy-preserving algorithms to encourage the stakeholders in sharing valuable insights without sacrificing their privacy concerns. We need to engage researchers and practitioners alike to identify more use cases for cooperation besides DDoS detection and mitigation.

Regulation and quality control of IoT devices

The design and deployment of IoT devices require a comprehensive overhaul to address bugs, misconfigurations, insecure defaults, and to incorporate support for patching and updates. But the intense competition among IoT device manufacturers to outdo one another in terms of feature support, and the benefits of being the "first to market" obliterate any incentives for manufacturers to invest more time and capital in quality control. Until recently there were no regulations governing the design and deployment of IoT devices. Even with the recent proposals [10, 44], the lack of a regulation-enforcing agency does not bode well for ensuring high quality of IoT devices.

4 SYSTEM DESIGN

We briefly discuss a broad spectrum of solutions highlighting, in particular, how each one falls short in addressing the DDoS security threats. We then present a high-level solution architecture and share key design tradeoffs that we must address in instantiating the solution.

4.1 Solution space

We classify the existing solutions to detect and mitigate large-scale DDoS attacks along two orthogonal dimensions. *First*, depending on the deployment location (e.g., access networks, and ISPs) the granularity level of the insights provided by a *security monitor*[3] varies. At the finest level the security monitor offers an *individual view*, allowing each IoT device to be independently tracked, and at the coarsest level it offers a *global view*, providing an aggregated view of the network activity. *Second*, a security monitor may track, and police the behavior of IoT devices either by presenting an *IP-level view* (finest granularity) or a *service-level view* (coarsest

[2] We note, for instance, that BGP owes its success as the de facto inter-domain routing protocol partly to its "information hiding" capabilities.
[3] We use the term "security monitor" to refer to a specific instantiation of a security framework.

granularity). We require insights from all four quadrants of the solution space, in Figure 1, to construct a practical, scalable security implementation.

Edge solutions This scenario entails deploying a security monitor either within an end user's home network, or inside the ISP or access network. Security monitors in a home network provide an individual and IP-level view, allowing fine-grained tracking of each IoT device. This approach, however, lacks the visibility required to detect a DDoS attack; the security monitor, however, makes it easier and safer to "quarantine" malicious IoT devices [36]. Deployments inside an ISP offer a more coarse-grained view, but the aggregated network activity from IoT devices in different home networks provide more insights into a network-wide DDoS attack. An ISP's monitor requires, however, cooperation of home networks to enforce a security policy, e.g., to precisely isolate one or more IoT devices from the Internet.

Core solutions In this case the security monitor is deployed either in a CDN or an Internet exchange point (IXP). CDNs have a widespread geographic footprint, and an extensive visibility into the Internet's traffic to observe, for instance, what services are being accessed by which IPs. Although CDNs can exploit their distributed infrastructure to monitor network traffic, centralized approaches to "scrub" traffic quickly render DDoS detection and mitigation economically unviable [31]. Similar to CDNs, IXPs allow monitoring of inter-network activity, but at a much coarser level (e.g., subnets). While it is also difficult to obtain deep service-level insights in IXPs, they offer a truly global view at a centralized location (i.e., where the IXP's switching fabric is located) making large-scale DDoS attack detection easier. It might be more useful, however, to detect a botnet during the early stages of its growth (i.e., when it has control of fewer devices) and subvert it completely rather than detect DDoS attacks after the botnet has reached its peak growth. The latter approach might often be too late for controlling large-scale attacks in practice.

4.2 A distributed, cooperative approach

Across these solutions there is a common thread that underscores the need for *cooperation* between the different entities of the Internet, e.g., IoT devices, home networks, ISPs, IXPs, and CDNs. To this end, we sketch a highly distributed (as in [23]), but concerted approach between the stakeholders, and outline the tradeoffs that we must address in instantiating the solution. Figure 2 illustrates the components of one instance of our proposed security monitor; our approach entails deploying several instances of this monitor at different locations, with the goal of gathering insights from different vantage points in a systematic manner to detect, mitigate, and control botnet attacks.

Data capture and analysis At the lowest-level of the implementation, the *data capture and handling* component comprises of modules to capture, process, and persist network activity data. This component parses the raw network traffic data, and converts it into a format that is easier to analyze. As part of this process, the component also discards redundant or irrelevant information to minimize the data storage overheads. The *data analysis and processing* component hosts a suite of different signal-processing algorithms to analyze the network data, and extract key attributes and insights from this data.

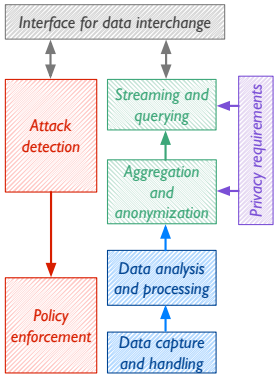

Figure 2: *Components in an instance of a security monitor implementation*

Data aggregation and sharing The *aggregation and anonymization* component is crucial for offering a scalable, and practical security monitor architecture. A security monitor can get quickly overwhelmed by the volume of traffic data analyzed (e.g., a security monitor instance deployed in an ISP or CDN). Data aggregation is vital to reduce the volume of data analyzed without losing critical insights. A smaller data footprint also reduces overheads associated with data sharing between instances and allows such cooperation to remain feasible in practice. Cooperation between the stakeholders, or the security monitor instances at different vantage points, is crucial for detecting and mitigating large-scale botnet attacks (refer §3). The aggregated data might contain sensitive information, and hence data anonymization is key to enable the exchange of data between stakeholders. Since requirements for privacy concerns vary among stakeholders, a separate *privacy requirements* module allows each stakeholder to explicitly declare their privacy requirements to security monitor instances deployed in their premises. The anonymization module uses the privacy requirements declared to tailor the anonymization algorithms as required. Lastly, the *streaming and querying* component supports streaming of analyses to other monitor instances, e.g., to alert another instance of a newly detected attack, and also querying of other instances for specific pieces of information, e.g., to gather more information on an ongoing attack.

Detection and enforcement The *attack detection* component hosts algorithms to detect different kinds of DDoS attacks. The attack detection algorithms may simply use only the locally available information, or may request additional data from other instances to confirm an attack signature. This component works in tandem with the streaming and querying component to accurately detect large-scale DDoS attacks. Detecting DDoS attacks is not of much significance if we cannot dissipate or thwart the attack to protect the attack target or victim. To this end, the *policy enforcement* component allows monitor instances to interact and enforce a security policy, e.g., remove the offending IoT devices from the Internet. While isolating the offending IoT devices is crucial to stop the attack, reaching hundreds of thousands of devices and severing their network connectivity takes time. In the meantime, the attackers

may take control of other vulnerable IoT devices. Hence, depending on the insights shared by the attack detection module, the policy enforcement module in each monitor crafts rules that are appropriate for that instance. Enforcement might entail a spectrum of different techniques from changing how traffic from a specific prefix or subnet is routed at an ISP, to dropping or "scrubbing" traffic destined towards a particular network or service at a CDN, to simply severing the network connectivity of specific IoT devices in home networks.

In summary, our approach is to combine the insights from multiple deployment locations, and exploit the Internet's hierarchical topology to our advantage. For instance, rather than detect an attack from a (logically) single vantage point, e.g., CDN, which might often prove too late to control or dissipate the attack, we leverage security monitors in home networks and ISPs, in addition to CDNs; while the monitors in home networks will only see a trickle of data from each device, the ISP may observe a pattern of traffic emanating from such devices towards one or more targets (or victims). Leveraging the topology offers a chance to detect botnets in their bootstrap phase and subvert them quickly before the attackers have any chance of exploiting the botnet. The location of the monitor in the topology also influences data aggregation and anonymization functions. A monitor in a home network is more concerned with protecting the identity of the users or devices in the network and has little to do in terms of aggregation. In contrast, a monitor in a ISP is concerned about both aggregating the data into meaningful prefixes, and protecting the identity of its users.

4.3 Implementation guidelines

In this section we present the guidelines for instantiating our IoT security architecture.

Interfaces for data interchange The interface for data interchange enables the stakeholders to exchange valuable insights, and, hence, cooperate with one another to detect and mitigate DDoS attacks. To this end, we present a couple of recommendations for designing the interface.

∗ While a home gateway router recognizes individual devices, a CDN's notion of a "client" or "user" rests at the IP level, and an IXP or ISP might only be interested at the level of a subnet. Naturally, the fine-grained insights of a security monitor in a home network should be transformed into more coarse-grained view to be useful to a monitor running in an ISP. Interfaces should be designed to facilitate such data aggregation as well as disaggregation when data is exchanged in the opposite direction, e.g., when an ISP shares coarse-grained insights with a monitor running in a home network.

∗ It is as important to attest that the responses obtained from the interfaces, e.g., traffic volume of IoT devices in a home network, are correct as it is to verify that the request is legitimate. Barring these mechanisms, an attacker can either fake results to masquerade an attack or trick the ISP into filtering out regular traffic, resulting in service disruptions. Lack of request authentication mechanisms could be exploited by attackers to gather sensitive data or inundate the security monitor by simply flooding illegitimate requests.

Scalable aggregation Exploiting the *differential observability* of monitor instances is crucial for DDoS attack detection. While it is impractical to send all data (i.e., insights) from an CDN's monitor instance to that of an IXP, naïve batching and aggregation can introduce delays in detection or even introduce errors. Besides, aggregation is often a resource-intensive task. There exist, hence, tradeoffs between bandwidth consumed and detection accuracy, and between detection delays and querying (or data sharing) frequency. Recent work on heavy-hitter detection [38] and distributed data processing [25] may, however, help in balancing these tradeoffs.

Privacy implications The exchange of data between monitor instances also has serious privacy implications for the involved parties. End users might object to unrestricted sharing of information on their home network to an ISP; the possibility of an ISP abusing the data shared for marketing or advertising, or, even worse, of an attacker stealing sensitive data about users calls for the use of privacy-preserving data-processing techniques, e.g., [11, 22]. A couple of promising approaches include *(a)* designing coarse-grained monitors, e.g., located in an ISP, to query fine-grained monitors, e.g., in home networks, and restrict the query processing or data analysis to *cloudlets* [15], and *(b)* restricting data analytics to trusted middleboxes that provide selective access without sacrificing the privacy requirements of stakeholders [41].

Real-time monitoring It is imperative to perform real-time monitoring of the growth of a botnet and detect DDoS attacks as early as possible, for instance, during the botnet's bootstrap phase. Large-scale DDoS attacks are difficult to dissipate or to control after the attack traffic exceeds a significant threshold, e.g., 100 Gbps of traffic volume. Although attackers can slow down the growth rate of a botnet, or shape the attack as required to escape detection, a concerted effort between different home networks can reveal the growth of a botnet, and help in subverting the attacks well before the attacker has even setup the DDoS infrastructure.

Policy enforcement Even if we quickly detect the next large-scale botnet attack, security enforcement opens a set of new challenges. Enforcing a security policy, e.g., removing the affected IoT devices from different home networks, is a hard problem. We need to engage the stakeholders in a debate to determine *(a)* who has the authority to deactivate "malicious" IoT devices, *(b)* how to turn off or sever the network connectivity of such IoT devices in user premises, and *(c)* the authentication mechanisms to allow legitimate deactivation requests. Incorrect policy enforcement can affect end-users' quality of experiences, e.g., by accidentally deactivating a user's Amazon Dash button, or even potentially endanger the end users, e.g., by incorrectly deactivating a user's home security system.

5 RELATED WORK

The unprecedented growth in IoT devices in the last few years combined with a shocking lapse of security mechanisms in these devices have resulted in an ecosystem that is ripe for exploitation by hackers. The fragility of the IoT ecosystem and its security implications, however, have been well-known to both academia and the industry [7, 8, 16, 33]. This paper complements these prior work by emphasizing why current approaches do not suffice to address the IoT security concerns.

Prior work have highlighted the lack of update mechanisms [9, 34, 40], poor design decisions [21, 29], and rampant use of default

credentials [13, 14] to quantify the fragility of the IoT ecosystem. Our objective in this position paper is to reignite this debate by highlighting the dire need for a coordinated effort among all stakeholders to secure the Internet from massive botnet attacks in the future.

Compared to prior work that propose point solutions, e.g., security managers within home routers [36], distributed data centers for "scrubbing" traffic [1], and attack dissipation strategies [30], we highlight the shortcomings of point solutions and call for a more coordinated effort among the stakeholders. Prior work addressing the privacy concerns [4, 15, 37] that naturally arise in the IoT ecosystem, or focussing on privacy-preserving analytics [11, 22] are very relevant to this position paper: they quell the privacy concerns that often stall cooperation between stakeholders.

6 CONCLUSION

The fragility of the Internet of Things (IoT) ecosystem endangers the security of the entire Internet. Due to the pervasive deployment of IoT devices in the Internet, we strongly caution the community to adopt a highly distributed, coordinated approach to address the security threats posed by the fragile IoT ecosystem. We highlight the inherent challenges in improving IoT security, and present guidelines and tradeoffs that researchers and practitioners must consider when instantiating an IoT security framework. While we do not solve the IoT security issues, we offer insights to inform the IoT security debate and bootstrap a discussion towards designing practical and scalable security solutions.

REFERENCES

[1] Akamai Technologies. 2014. Prolexic Routed: DDoS defense for protecting data center infrastructures against large, complex attacks. "https://goo.gl/eAdko6". (October 2014).

[2] Akamai Technologies. 2016. State of the Internet Security Spotlight, Q4 2016: Internet of Things and the Rise of 300 Gbps DDoS Attacks. "https://goo.gl/1uHJuY". (2016).

[3] Akamai Technologies. 2017. State of the Internet / Security, Q1 2017 Report. "https://goo.gl/KJ4oNX". (May 2017).

[4] Tristan Allard, Davide Frey, George Giakkoupis, and Julien Lepiller. 2016. Lightweight Privacy-Preserving Averaging for the Internet of Things. In *Proceedings of the 3rd Workshop on Middleware for Context-Aware Applications in the IoT (M4IoT 2016)*. ACM, New York, NY, USA, 19–22.

[5] Manos Antonakakis, Tim April, Michael Bailey, Matt Bernhard, Elie Bursztein, Jaime Cochran, Zakir Durumeric, J. Alex Halderman, Luca Invernizzi, Michalis Kallitsis, Deepak Kumar, Chaz Lever, Zane Ma, Joshua Mason, Damian Menscher, Chad Seaman, Nick Sullivan, Kurt Thomas, and Yi Zhou. 2017. Understanding the Mirai Botnet. In *26th USENIX Security Symposium (USENIX Security 17)*.

[6] Chris Baker. 2016. Recent IoT-based Attacks: What Is the Impact On Managed DNS Operators? "https://goo.gl/kCU7ZL". (October 2016).

[7] Mario Ballano Barcena and Candid Wueest. 2015. *Insecurity in the Internet of Things*. Technical Report. Symantec.

[8] E. Bertino and N. Islam. 2017. Botnets and Internet of Things Security. *Computer* 50, 2 (Feb 2017), 76–79.

[9] H Birkholz, N Cam-Winget, and C Bormann. 2016. IoT Software Updates need Security Automation. *Internet of Things Software Update Workshop (IoTSU)* (May 2016).

[10] Matthew Broersma. 2016. EU Pushes IoT Security Regulations. "https://goo.gl/CiTA9P". (October 2016).

[11] Ruichuan Chen, Alexey Reznichenko, Paul Francis, and Johanes Gehrke. 2012. Towards Statistical Queries over Distributed Private User Data. In *Presented as part of the 9th USENIX Symposium on Networked Systems Design and Implementation (NSDI 12)*. USENIX, San Jose, CA, 169–182.

[12] Alex Chiu. 2016. The Internet of Things Is Not Always So Comforting. "https://goo.gl/WxELos". (February 2016).

[13] Catalin Cimpanu. 2017. 15% of All IoT Device Owners Don't Change Default Passwords. "https://goo.gl/H2RLHa". (June 2017).

[14] Ang Cui and Salvatore J. Stolfo. 2010. A Quantitative Analysis of the Insecurity of Embedded Network Devices: Results of a Wide-area Scan. In *Proceedings of the 26th Annual Computer Security Applications Conference (ACSAC '10)*. ACM, New York, NY, USA, 97–106.

[15] Nigel Davies, Nina Taft, Mahadev Satyanarayanan, Sarah Clinch, and Brandon Amos. 2016. Privacy Mediators: Helping IoT Cross the Chasm. In *Hot Topics in Mobile Computing (Hot Mobile)*.

[16] Tamara Denning, Tadayoshi Kohno, and Henry M. Levy. 2013. Computer Security and the Modern Home. *Commun. ACM* 56, 1 (Jan. 2013), 94–103.

[17] Dave Evans. 2011. *The Internet of Things: How the Next Evolution of the Internet Is Changing Everything*. Technical Report. CISCO.

[18] Dennis Fisher. 2016. Bugs in Chinese IoT Components Aid Mirai Botnet Spread. "https://goo.gl/iGcwdt". (October 2016).

[19] Gartner Inc. 2017. Gartner Says 8.4 Billion Connected "Things" Will Be in Use in 2017, Up 31 Percent From 2016. "https://goo.gl/JVRWXV". (February 2017).

[20] Imperva. 2017. *Attackers Use DDoS Pulses to Pin Down Multiple Targets, Send Shock Waves Through Hybrids*. Technical Report. Imperva Incapsula.

[21] Brian Krebs. 2016. IoT Reality: Smart Devices, Dumb Defaults. "https://goo.gl/PE3iif". (February 2016).

[22] D. Le Quoc, M. Beck, P. Bhatotia, R. Chen, C. Fetzer, and T. Strufe. 2017. Privacy Preserving Stream Analytics: The Marriage of Randomized Response and Approximate Computing. *ArXiv e-prints* (Jan. 2017).

[23] Tom Leighton. 2009. Improving Performance on the Internet. *Commun. ACM* 52, 2 (Feb. 2009), 44–51.

[24] Robert McMillan and Drew FitzGerald. 2016. Hackers Release Botnet Code, Raising Specter of More Attacks. "https://goo.gl/RjsMdd". (October 2016).

[25] Chris Olston, Jing Jiang, and Jennifer Widom. 2003. Adaptive Filters for Continuous Queries over Distributed Data Streams. In *Proceedings of the 2003 ACM SIGMOD International Conference on Management of Data (SIGMOD '03)*. ACM, New York, NY, USA, 563–574.

[26] Martin Orehek and Alf Zugenmaier. 2016. Updates in IoT are more than just one iota. *Internet of Things Software Update Workshop (IoTSU)* (2016).

[27] Pierluigi Paganini. 2016. OVH hosting hit by 1Tbps DDoS attack, the largest one ever seen. "https://goo.gl/2raB5o". (September 2016).

[28] Andrea Peterson. 2016. Can anyone keep us safe from a weaponized 'Internet of Things?'. "https://goo.gl/yzbzrx". (October 2016).

[29] David Plonka and Elisa Boschi. 2016. The Internet of Things Old and Unmanaged. *Internet of Things Software Update Workshop (IoTSU)* (2016).

[30] Matthew Prince. 2016. How Cloudflare's Architecture Allows Us to Scale to Stop the Largest Attacks. "https://goo.gl/xh4J8j". (October 2016).

[31] Steve Ragan. 2016. Some thoughts on the Krebs situation: Akamai made a painful business call. "https://goo.gl/PVdz3F". (September 2016).

[32] David E. Sanger and Nicole Perlroth. 2016. A New Era of Internet Attacks Powered by Everyday Devices. "https://goo.gl/djNdTC". (October 2016).

[33] Bruce Schneier. 2014. The Internet of Things is Wildly Insecure — and often Unpatchable. "https://goo.gl/CG87c9". (January 2014).

[34] Bruce Schneier. 2017. Ransomware and the Internet of Things. "https://goo.gl/ZC2vMF". (May 2017).

[35] Chad Seaman. 2016. *Akamai Threat Advisory: Mirai Botnet*. Technical Report. Akamai Technologies.

[36] A. K. Simpson, F. Roesner, and T. Kohno. 2017. Securing vulnerable home IoT devices with an in-hub security manager. In *2017 IEEE International Conference on Pervasive Computing and Communications Workshops (PerCom Workshops)*. 551–556.

[37] Rayman Preet Singh, Benjamin Cassell, S. Keshav, and Tim Brecht. 2016. TussleOS: Managing privacy versus functionality trade-offs on IoT devices. In *Computer Communication Review (CCR)*, Vol. 46. ACM SIGCOMM, 1–8.

[38] Vibhaalakshmi Sivaraman, Srinivas Narayana, Ori Rottenstreich, S. Muthukrishnan, and Jennifer Rexford. 2017. Heavy-Hitter Detection Entirely in the Data Plane. In *Proceedings of the Symposium on SDN Research (SOSR '17)*. ACM, New York, NY, USA, 164–176.

[39] Robin Sommer and Vern Paxson. 2010. Outside the Closed World: On Using Machine Learning for Network Intrusion Detection. In *Proceedings of the 2010 IEEE Symposium on Security and Privacy (SP '10)*. IEEE Computer Society, Washington, DC, USA, 305–316.

[40] Robert Sparks and Ben Campbell. 2016. Avoiding the Obsolete-Thing Event Horizon. *Internet of Things Software Update Workshop (IoTSU)* (2016).

[41] Nik Sultana, Markulf Kohlweiss, and Andrew W. Moore. 2016. Light at the Middle of the Tunnel: Middleboxes for Selective Disclosure of Network Monitoring to Distrusted Parties. In *Proceedings of the 2016 Workshop on Hot Topics in Middleboxes and Network Function Virtualization (HotMIddlebox '16)*. ACM, New York, NY, USA, 1–6.

[42] US-CERT. 2016. Alert (TA16-288A): Heightened DDoS Threat Posed by Mirai and Other Botnets. "https://goo.gl/SYA8JW". (November 2016).

[43] Jai Vijayan. 2017. 'Pulse Wave' DDoS Attacks Emerge As New Threat. "https://goo.gl/aw8g78". (August 2017).

[44] Mark R. Warner, Cory Gardner, Ron Wyden, and Steve Daines. 2017. Internet of Things Cybersecurity Improvement Act of 2017. "https://goo.gl/DZESrV". (2017).

[45] Nick Woolf. 2016. TheGuardian: DDoS attack that disrupted internet was largest of its kind in history, experts say. "https://goo.gl/HWsNyx". (October 2016).