

# MITIGATING THE INCREASING RISKS OF AN INSECURE INTERNET OF THINGS

NICK FEAMSTER\*

*The Internet of Things (IoT) comprises Internet-connected devices that serve a special function—ranging from personal health devices to environmental sensors—in contrast to general-purpose computing devices, such as laptops or smartphones. The emergence and proliferation of IoT devices on industrial, enterprise, and home networks brings with it unprecedented risk. The potential magnitude of this risk was made concrete in October 2016, when insecure Internet-connected cameras launched a distributed denial of service (DDoS) attack on Dyn, a provider of Domain Name System (DNS) service for many large online service providers (e.g., Twitter, Reddit). Although this incident caused large-scale disruption, it is noteworthy that the attack involved only a few hundred thousand endpoints and a traffic rate of about 1.2 terabits per second.<sup>1</sup> With predictions of upwards of a billion new IoT devices within the next five to ten years,<sup>2</sup> the risk of similar, and likely much larger, attacks is imminent.*

*In this Article, I provide an overview of the growing risks associated with insecure connected IoT devices and present various approaches that may ultimately help mitigate these risks. Many of the approaches that this Article posits depend on technical solutions that are as of yet incomplete; in some cases, they also depend on the alignment of incentives between various stakeholders in the IoT ecosystem. In this sense, this Article offers more questions than it answers; my aim is to point the community in potentially fruitful directions for studying technical and regulatory-related questions concerning IoT security.*

---

\* Nick Feamster is a Professor of Computer Science at Princeton University's Center for Information Technology Policy.

1. Nicky Woolf, *DDoS Attack that Disrupted Internet was Largest of its Kind in History, Experts Say*, GUARDIAN (Oct. 26, 2016, 4:42 PM), <https://www.theguardian.com/technology/2016/oct/26/ddos-attack-dyn-mirai-botnet> [<https://perma.cc/26XU-6XBU>].

2. JAMES MANYIKA ET AL., THE INTERNET OF THINGS: MAPPING THE VALUE BEYOND THE HYPE (2015).

INTRODUCTION .....	88
I. ENCOURAGING STAKEHOLDERS TO MITIGATE AND MANAGE RISK .....	90
A. <i>Device Manufacturers</i> .....	90
B. <i>Consumers</i> .....	93
C. <i>Internet Service Providers (ISPs)</i> .....	95
D. <i>Summary</i> .....	96
II. IMPROVING INFRASTRUCTURE RESILIENCE .....	97
A. <i>Device Identification and Inventory</i> .....	98
B. <i>Visibility and Control over Traffic Flows</i> .....	99
C. <i>Improving the Resilience of the Internet Infrastructure</i> .....	100
CONCLUSION .....	101

## INTRODUCTION

Although connected devices generally pose security risks for the Internet at large, the scale and scope of the Internet of Things (IoT) presents unprecedented risks to Internet security. IoT presents unique security challenges due to the sheer number of IoT devices that users are connecting to the Internet, with predictions of more than a billion connected IoT devices in the next few years. The heterogeneity of these devices and the manufacturers who make them also poses unprecedented risks; each device may behave differently, making it difficult to establish a baseline for “normal” traffic. Furthermore, even if it were easy to detect device misbehavior, the large expanse of IoT device manufacturers makes it more difficult to hold manufacturers accountable when their devices introduce security risks.

One of the biggest contributors to the risk of future attacks is the fact that many IoT devices have long-standing, widely-known software vulnerabilities that make them prone to exploit and control by remote attackers.<sup>3</sup> Worse yet, the vendors of these IoT devices often have provenance in the hardware industry, but they may lack expertise or resources in software development and systems security. As a result, IoT device manufacturers may ship devices that are extremely difficult, if not practically impossible, to secure. The large number of insecure IoT devices connected to the Internet poses unprecedented risks to consumer privacy, as well as threats to the underlying physical infrastructure and the global Internet at large:

---

3. Brian Krebs, *IoT Reality: Smart Devices, Dumb Defaults*, KREBS ON SECURITY (Feb. 8, 2016, 10:15 AM), <http://krebsonsecurity.com/2016/02/iot-reality-smart-devices-dumb-defaults/> [https://perma.cc/MBH8-YJU8].

- **Data privacy risks.** Internet-connected devices increasingly collect data about the physical world, including information about the functioning of infrastructure, such as the power grid and transportation systems, as well as personal or private data on individual consumers. At present, many IoT devices either do not encrypt their communications or use a form of encrypted transport that is vulnerable to attack.<sup>4</sup> Many of these devices also store the data they collect in cloud-hosted services, which may be the target of data breaches or other attack.
- **Risks to availability of critical infrastructure and the Internet at large.** As the Mirai botnet attack of October 2016 demonstrated, Internet services often share core dependencies on the underlying infrastructure. Crippling many websites offline did not require direct attacks on these services, but rather a targeted attack on the underlying infrastructure on which many of these services depend (i.e., the Domain Name System). More broadly, one might expect future attacks that target not just the Internet infrastructure, but also physical infrastructure that is increasingly Internet-connected (e.g., power and water systems). The dependencies that are inherent in the current Internet architecture create immediate threats to resilience.

The large magnitude and broad scope of these risks implore us to seek solutions that will improve infrastructure resilience in the face of Internet-connected devices that are extremely difficult to secure. A central question in this problem area concerns the responsibility that each stakeholder in this ecosystem should bear, and the respective roles of technology and regulation (whether via industry self-regulation or otherwise) in securing both the Internet and associated physical infrastructure against these increased risks.

In this Article, I will discuss various approaches toward reducing the increasing risks of an insecure IoT. In the first part of the Article, I will enumerate the various stakeholders and the various technical and policy levers that might be used to encourage risk-mitigating behavior from each of the corresponding stakeholders. I will then suggest various technical approaches for managing and mitigating the risks of insecure connected IoT devices, along with the associated challenges of implementing these approaches in practice.

---

4. Charlie Osborne, *Internet of Things devices lack fundamental security, study finds*, ZDNET (Apr. 8, 2015, 1:45 PM), <http://www.zdnet.com/article/internet-of-things-devices-lack-fundamental-security-study-finds/> [<https://perma.cc/G7MU-YXZT>].

## I. ENCOURAGING STAKEHOLDERS TO MITIGATE AND MANAGE RISK

The three main stakeholders in the IoT ecosystem are: (1) the device manufacturers, who must make decisions about the extent to which they incorporate security best practices into their products; (2) the consumers, who must make decisions about both the products that they buy and the extent to which they apply software updates to these products; and (3) the Internet service providers (ISPs), who can potentially take action against consumers or manufacturers who do not abide by the practice of “good security hygiene.” In this section, I explore various roles that these stakeholders might play in securing the IoT ecosystem, as well as various challenges in implementing these approaches.

### A. *Device Manufacturers*

One possible lever for either government or self-regulation is the IoT device manufacturers. Device manufacturers could potentially be held accountable by regulators (potentially through fines), through ISPs (who could take steps to limit if and how these devices may communicate with other destinations on the Internet), and consumers (who make purchasing decisions and could also install technology in their homes that might allow them to better expose and control device behavior).

One possibility might be a device certification program for manufacturers that could attest to an adherence to best common practice for device and software security. A well-known (and oft-used) analogy is the Underwriters Laboratories (UL) certification process for electrical devices and appliances.<sup>5</sup> A certification process typically checks that a device meets some type of manufacturing standard. In the case of UL certification, a device bearing a UL mark would indicate that UL has tested instances of a product and has determined that those instances meet the security requirements that are based on accepted standards. In the case of IoT, such a certification standard might include specifications about software update processes, communication restrictions, or approaches to storing sensitive user data.

Despite its conceptual appeal, however, a certification approach poses several practical challenges. One challenge is outlining and prescribing best common practices in the first place, particularly due to the rate at which technology (and attacks) progress. Any specific set of prescriptions runs the risk of falling out of date as technology advances. Similarly, certification can readily devolve into a checklist of attributes that vendors satisfy, without necessarily adhering to the

---

5. Specific Guidelines and Rules, UL, <http://www.ul.com/marks/ul-listing-and-classification-marks/promotion-and-advertising-guidelines/specific-guidelines-and-rules/> [https://perma.cc/AGZ2-3JX7] (last visited Oct. 16, 2017).

*process* by which these devices are secured over time. For example, certification could specify a set of protocols or processes that a device must incorporate (e.g., a specific version of the Transport Layer Security protocol,<sup>6</sup> or a specific implementation of the protocol in a version of the OpenSSL library), yet, over time, specific versions of protocols and the libraries that implement them may be discovered to be insecure.

Although certification should not be overly prescriptive, recent reports, such as that from the Broadband Internet Technology Advisory Group (BITAG), outline specific practices that are not specific to any particular protocol or technology, but nonetheless highlight a set of best practices that manufacturers should abide by. These recommendations include the following:

- **Devices should use best current software practices.** Devices should ship with reasonably current software, incorporate a mechanism for automated, secure software updates, use strong authentication by default, and use configurations that have undergone extensive testing and hardening.
- **Devices should follow best practices with respect to security and cryptography.** Devices should encrypt communications—both with IoT controllers and to cloud servers—by default; encrypt local storage; authenticate communications, software changes, and requests for data; close unnecessary ports and disable unnecessary services; and use libraries that are actively maintained and supported. Devices should also be restrictive, rather than permissive, in communicating, meaning that they should not rely solely on network firewalls to restrict their communications, but should, in general, be restrictive about the devices that they engage in communications with.
- **Device function should be robust to disruptions to connectivity for the availability of cloud-back-end services.** To the extent possible, devices that depend on a cloud back-end service or Internet connectivity should continue to provide basic function, even if connectivity is disrupted or the cloud back-end service is disrupted or fails. For example, a smart light switch or thermostat should continue to function as a switch or thermostat, respectively, even when Internet connectivity is lost or the supporting cloud service fails.

---

6. T. DIERKS & E. RESCORLA, THE TRANSPORT LAYER SECURITY (TLS) PROTOCOL (2008), <https://tools.ietf.org/html/rfc5246> [<https://perma.cc/QUY2-NLSW>].

The complete BITAG report contains additional recommendations.<sup>7</sup>

As daunting as challenges of specifying a certification program may seem, encouraging adherence to a certification program may prove even more challenging. This concern may be particularly acute for consumer IoT, where consumers may not bear the direct costs of connecting insecure devices to their home networks. Specifically, consumers may not appreciate the value of certification, particularly if meeting the requirements of certification increases the cost of a device. Ultimately, as I discuss in the subsequent section, the costs of connecting devices to the network that do not meet minimum specification may need to be passed to consumers to encourage them to purchase specific devices.

Device distribution channels (i.e., retailers) may ultimately play an important role in communicating the practices of device manufacturers to consumers and ultimately encouraging consumers to opt for more secure IoT devices. For example, given the existence of a certification program, a distributor or retailer could indicate whether a particular device passed some level of certification. An online retailer could display a device's certification status prominently, and it could also preferentially order a user's search results according to which devices are certified (e.g., by listing certified devices at the top of a list of search results).

Another challenge with device certification is providing convenient mechanisms to device manufacturers to *implement* the recommendations prescribed by a certification program. Ultimately, device manufacturers face a potential cost when implementing a set of recommendations, so lowering the cost of implementing the recommendations is an important consideration. One way to lower these costs is to develop a software library implementing these recommendations that is easy for device manufacturers to incorporate into the systems that they are developing. Ongoing work from IoTivity<sup>8</sup> and the Open Connectivity Foundation (OCF)<sup>9</sup> is in the process of creating such libraries and is a promising development in this regard. As open-source software libraries become available for developing IoT applications and services, these libraries should lower the barrier for deploying IoT devices that conform to the specifications outlined by a certification program.

---

7. BROADBAND INTERNET TECH. ADVISORY GRP., INTERNET OF THINGS (IoT) SECURITY AND PRIVACY RECOMMENDATIONS (2016), [https://www.bitag.org/documents/BITAG\\_Report\\_-\\_Internet\\_of\\_Things\\_\(IoT\)\\_Security\\_and\\_Privacy\\_Recommendations.pdf](https://www.bitag.org/documents/BITAG_Report_-_Internet_of_Things_(IoT)_Security_and_Privacy_Recommendations.pdf) [https://perma.cc/3WBT-2NUB].

8. IOTIVITY, <https://www.iotivity.org/> [https://perma.cc/JQ3U-XYBS] (last visited Oct. 16, 2017).

9. OPEN CONNECTIVITY FOUND., <https://openconnectivity.org/> [https://perma.cc/8SZD-HZWZ] (last visited Oct. 16, 2017).

## B. Consumers

Consumers are another group of stakeholders who could be incentivized to improve the security of the devices that they connect to their networks (in addition to more effectively securing the networks to which they connect these devices). As the entity who purchases and ultimately connects IoT devices to the network, consumers appear well-situated to ensure the security of the IoT devices on their respective networks.

Unfortunately, the picture is a bit more nuanced. First, consumers typically lack either the aptitude or interest (or both) to secure either their own networks or the devices that they connect to them.<sup>10</sup> For example, home broadband Internet access users have generally proved to be poor at applying software updates in a timely fashion, and have been equally delinquent in securing their home networks.<sup>11</sup> Even skilled network administrators regularly face network misconfigurations, attacks, and data breaches.<sup>12</sup> Ongoing surveys of users of smart-home equipment indicate that users are often unaware of the security and privacy risks associated with the devices that they are deploying in their networks. Consumers often blindly trust manufacturers, assuming them to be responsible purveyors of their data, despite the fact that devices have repeatedly proven to be insecure.

Second, in many cases, users may lack the incentives to ensure that their devices are secure. In the case of the Mirai botnet, for example, consumers did not directly face the brunt of the attack; rather, the ultimate victims of the attack were DNS service providers and, indirectly, online service providers, such as Twitter.<sup>13</sup> To the first order, consumers suffered little direct consequence as a result of insecure devices on their networks. The basic problem relates to externalities: users face additional inconvenience and cost from securing their devices, but their failure to do so does not typically result in direct negative outcomes, which are more often faced by other Internet services and users. These problems are reminiscent of well-studied problems in the economics of information security, and some of the techniques that have applied in past settings may also be applicable to IoT.

---

10. Rebecca E. Grinter et al., *The Work to Make a Home Network Work*, in PROCEEDINGS OF THE NINTH EUROPEAN CONFERENCE ON COMPUTER-SUPPORTED COOPERATIVE WORK (2005).

11. Arunesh Mathur et al., *“They Keep Coming Back Like Zombies”: Improving Software Updating Interfaces*, in USENIX, PROCEEDINGS OF THE TWELFTH SYMPOSIUM ON USABLE PRIVACY AND SECURITY (2016), <https://www.usenix.org/system/files/conference/soups2016/soups2016-paper-mathur.pdf> [<https://perma.cc/C5D5-PSWN>].

12. Ratul Mahajan et al., *Understanding BGP Misconfiguration*, 32 ACM SIGCOMM COMPUT. COMMUNIC’N REV. 3 (2002).

13. Scott Hilton, *Dyn Analysis Summary Of Friday October 21 Attack*, ORACLE DYN (Oct. 26, 2016), <https://dyn.com/blog/dyn-analysis-summary-of-friday-october-21-attack/> [<https://perma.cc/KYT8-76QX>].

Consumers' misaligned incentives suggest several possible courses of action. One approach might involve placing some responsibility or liability on consumers for the devices that they connect to the network, in the same way that a citizen might be fined for other transgressions that have externalities (e.g., fines for noise or environmental pollution). Alternatively, ISPs, or another entity, might offer users a credit for purchasing and connecting only devices that pass certification; another variation of this approach might require users to purchase "Internet insurance" from their ISPs that could help offset the cost of future attacks.<sup>14</sup> Consumers might receive credits or lower premiums based on the risk associated with their behavior (i.e., their software update practices or results from security audits of devices that they connect to the network).

In one possible scenario, an ISP might actively scan the network for vulnerable devices, using a suite of off-the-shelf vulnerability scanners, or perhaps a custom suite of tests that could test conformance to a set of best practices outlined in a certification program. Given the outcome of such a vulnerability scan, the ISP could ascertain a level of risk associated with a particular consumer.

Such a risk assessment could then be used to determine whether that consumer should shoulder some of the potentially increased costs that might result from an attack, such as those we have witnessed in the past year. This cost could be framed to the consumer in several different ways. One possible framing is through increased fees on a monthly bill. Given the relative unpopularity of ISPs, however, such a cost might better be framed as a "good hygiene discount" for users whose devices conform to certification, as opposed to a penalty levied on users whose devices do not conform. Another possible scenario might have consumers bear the costs of an attack, should their devices end up participating in or enabling such an attack. In such a scenario, consumers might have the option to purchase insurance against such attacks, where insurance premiums could be adjusted depending on the amount of risk that a consumer incurs by connecting a particular set of devices to the network.

In addition to monetary incentives, privacy might also prove to be a catalyst for consumer behavior. For example, IoT devices that are insecure may also often have poor privacy practices, either with respect to the device itself or with respect to the back-end services that store user data. Already, the IoT industry has seen countless cases of devices that fail to secure the communications between the device and the cloud, or that have been subject to data breaches or other incidents involving user data. Although users do not directly bear the costs of a denial of service attack, they may directly bear the

---

14. Jean Bolot & Marc Lelarge, *Cyber Insurance as an Incentive for Internet Security*, in *MANAGING INFORMATION RISK AND THE ECONOMICS OF SECURITY* 269–290 (M. Eric Johnson ed., 2008).



costs of a data breach. While it has not been determined that these events are correlated (i.e., that devices that mishandle user data are also more likely to participate in large-scale Internet attacks), it is possible that raising user awareness about privacy risks of IoT devices may spur consumers to purchase devices that have better privacy practices, which might indirectly mitigate other security risks. This possibility warrants further study.

### C. Internet Service Providers

A third stakeholder to consider is the ISP who provides Internet connectivity to the consumer. The ISP has considerable incentives to ensure that the devices that its customer connect to the network are secure: insecure devices increase the presence of attack traffic and may ultimately degrade Internet service or performance for the rest of the ISPs' customers.<sup>15</sup> From a technical perspective, the ISP is also in a uniquely effective position to detect and squelch attack traffic coming from IoT devices. The ISP also has unique visibility into a consumer's home network. By virtue of having a device in the home network, the ISP can often determine the devices that are connected to the Internet through the home network.<sup>16</sup> This device potentially allows the ISP to pinpoint the source of an attack or perhaps even stop the attack traffic entirely by firewalling traffic to or from the IoT device directly at the ISP-provided modem or access point (typically referred to as "customer premises equipment," or CPE).

Unfortunately, relying on the ISP alone to protect the network against insecure IoT devices is fraught with complications, some technical and others non-technical. One technical challenge concerns detecting anomalous IoT device traffic. Given traffic to or from an IoT device, an ISP could potentially determine whether the device participated in an attack, such as a denial of service attack. A challenge with anomaly detection in this context is that Internet traffic is increasingly becoming encrypted,<sup>17</sup> making it difficult to identify attack traffic using simple inspection techniques. Nevertheless, I believe that certain traffic features—such as the DNS domain names that a device looks up or the volume of traffic it is sending to individual destinations—may prove useful for identifying

---

15. David Moore et al., *Inferring Internet Denial-of-Service Activity*, 24 ACM TRANSACTIONS ON COMPUTER SYSTEMS 115 (2006).

16. Sarthak Grover et al., *Peeking Behind the NAT: An Empirical Study of Home Networks*, in PROCEEDINGS OF THE 2013 CONFERENCE ON INTERNET MEASUREMENT CONFERENCE (2013), <http://conferences.sigcomm.org/imc/2013/papers/imc061-groverA.pdf> [<https://perma.cc/SA/P5-8MBS>].

17. GLOBAL INTERNET PHENOMENA SPOTLIGHT: ENCRYPTED INTERNET TRAFFIC, SANDVINE, <https://www.sandvine.com/downloads/general/global-internet-phenomena/2015/encrypted-internet-traffic.pdf> [<https://perma.cc/3686-YARM>] (last visited Oct. 16, 2017).

anomalous behavior. I am exploring some of these possibilities in my ongoing research on Internet anomaly detection.

One non-technical challenge concerns the ISP's relationship with its customers: while the ISP could technically defend against an attack by disconnecting or firewalling consumer devices that are launching attacks, such an approach will certainly result in increased complaints and technical support calls from customers, who connect devices to the network and simply expect them to work. Better coordination with device manufacturers may ultimately enable ISPs to perform these types of firewalling actions. For example, as previously discussed, device manufacturers could ensure that IoT devices continued to perform basic functions even when not connected to the Internet; in this case, an ISP could identify that a device was misbehaving and firewall it, without preventing a customer from continuing to use the basic functions of the device.

A second challenge relates to privacy. Many of the technical capabilities that an ISP might have at its disposal (e.g., the ability to identify attack traffic coming from a specific device) introduce serious privacy concerns. For example, being able to alert a customer to, say, a compromised baby monitor requires the ISP to know (and document) that a consumer has such a device in the first place. These privacy challenges are fundamentally difficult for an ISP to overcome. On the one hand, vague alerts concerning the insecure behavior of a device in the home (e.g., "Your home network appears to have an insecure device.") are to nobody's benefit—they are not actionable for the user, and they are likely to generate increased complaints and tech support calls from customers. A potential solution might involve providing the consumer with additional tools to diagnose the source of the problem themselves (e.g., software running locally on the home network that exposes the behavior of individual devices), but doing so places additional onus on the consumer, which may prove impractical.

#### *D. Summary*

Managing the increased risks associated with insecure IoT devices will likely require action from all three stakeholders: device manufacturers, consumers, and ISPs. Some of the salient challenges will concern how the risks can be best balanced against the higher operational costs that will be associated with improving security, as well as who will ultimately bear these responsibilities and costs. A path forward almost certainly involves some form of regulation. This regulation might be driven by the market; as described above, vendors could participate in a certification program, which retailers could then use to signal to consumers, who would bear the benefits or costs of their purchasing decisions. Whether this type of structure

evolves naturally as self-regulation, or a regulatory agency ultimately will need to catalyze such a structure, remains to be seen.

Recent developments in the United States Government suggest some attention to IoT security, but we have yet to see a comprehensive approach. For example, Senator Warner recently proposed a bill that sets security standards for IoT devices in the context of government procurement. The Department of Commerce<sup>18</sup> and Federal Trade Commission<sup>19</sup> have also published reports indicating some desire to improve IoT security.

It is difficult to predict whether these initial developments suggest a more comprehensive or coherent regulatory approach. As with any standardization effort, the practical difficulties often concern implementation: Developers of software for IoT devices need convenient ways to integrate these approaches into their products. As such, a more likely a fruitful short-term outcome may be some type of market self-regulation, through organizations such as the Open Connectivity Foundation.

## II. IMPROVING INFRASTRUCTURE RESILIENCE

In addition to improving defenses against the insecure devices themselves, it is also critical to determine how to better build resilience into the underlying Internet infrastructure to cope with these attacks. If one views the occasional IoT-based attack as inevitable to some degree, one major concern is ensuring that the Internet infrastructure (and the associated cyberphysical infrastructure) remains both secure and available in the face of attack.

In the case of the Mirai attack on Dyn, for example, the severity of the attack was exacerbated by the fact that many online services depended on the infrastructure that was attacked. Computer scientists and Internet engineers should be thinking about technologies that can both potentially decouple these underlying dependencies and ensure that the infrastructure itself remains secure even in the event that regulatory or legal levers fail to prevent every attack. One possibility that I am exploring in my own research, for example, is the role that an automated home network firewall could play in (A) helping users keep better inventory of connected IoT devices, and (B) providing users both visibility into and control over the traffic flows that these devices send. I outline each of these possibilities in more detail below.

---

18. INTERNET POLICY TASK FORCE & DIGITAL ECON. TEAM, DEP'T OF COMMERCE, FOSTERING THE ADVANCEMENT OF THE INTERNET OF THINGS (2017).

19. FED. TRADE COMM'N, INTERNET OF THINGS: PRIVACY & SECURITY IN A CONNECTED WORLD (2015).

### A. *Device Identification and Inventory*

An important aspect of IoT device security is ensuring that the devices that are connected to the network can be enumerated and identified. Although this task appears simple at first blush, it is complicated by the fact that devices do not typically automatically identify themselves. As a result, device identification must rely on analysis of characteristics of the network traffic, such as the device's hardware address, as well as the DNS domain names that the device looks up (which can occasionally even uniquely identify the device). Enumerating and identifying devices is important for a couple of reasons:

- It allows network analysis tools and techniques to establish baseline behavior, which is often specific to the device.
- Given knowledge of the specific device that is exchanging traffic, network monitoring applications can provide users with actionable information about the source of offending traffic.

As mentioned, the two most common mechanisms for identifying devices are to use the device's hardware address (sometimes referred to as the "MAC address") and the domain names that the device looks up. A hardware address is a 48-bit address that is unique to the device; the first 24 bits of this address are reserved for a manufacturer identifier, which may sometimes identify the vendor for a particular device (e.g., Apple, Samsung). Unfortunately, this information typically does not permit device identification to a finer granularity (e.g., Samsung Galaxy S8), and because it refers to the *chipset* manufacturer, sometimes the manufacturer ID is even less helpful.

The domain names that a device looks up can also sometimes help identify a device. For example, we observed experimentally that the Nest Dropcam periodically issues DNS queries of the form `nexus.dropcam.com`.<sup>20</sup> In such cases, the DNS lookups that a device issues—which are always unencrypted—can identify a device. Sometimes, however, a device will issue queries to the domain name of a manufacturer (e.g., `nest.com`), but those lookups may not uniquely identify the type of device that issues the query. In these cases, the DNS traffic does not permit precise identification of the device, but may still help narrow down the device to a specific manufacturer or class of devices.

---

20. Nolan Apthorpe et al., *A Smart Home is No Castle: Privacy Vulnerabilities of Encrypted IoT Traffic*, <http://datworkshop.org/papers/dat16-final37.pdf> [https://perma.cc/PGW9-5F46] (last visited Nov. 3, 2017).

Both of these mechanisms may be further complicated by the introduction of IoT smart hubs, which connect a collection of devices in the home and originate traffic flows themselves. In these cases, all traffic from devices behind the hub are effectively “mixed,” making it appear as though the traffic is coming from a single device (i.e., the hub), as opposed to individual devices. Enumerating devices behind an IoT hub and disambiguating the source of each respective traffic flow remains an important, open technical challenge.

### B. *Visibility and Control over Traffic Flows*

Consumers need better *visibility* into the data flows that their devices send and *control* over which flows should be permitted or denied. In the common-case scenario in today’s home network, a user may connect an IoT device to the network yet have no insight into how that device is exchanging traffic with either other devices on the same home network or other general Internet destinations. Worse yet, users have little ability to control whether, how, and under what circumstances a device exchanges traffic with other network endpoints.

The Federal Trade Commission has recently noted these shortcomings and encouraged the broader community to provide solutions to these problems as part of its IoT Home Inspector Challenge competition.<sup>21</sup> Toward the goals of providing users better visibility into and control over traffic flows, my research lab has recently developed software that can run on a small device in the user’s home network (e.g., a Raspberry Pi), analyze the traffic that each device is sending, and potentially allow the user to block offending traffic flows. Such techniques rely first on the ability to attribute traffic flows to specific devices, a task that can prove difficult if the IoT devices are connected behind an IoT hub. These techniques also assume that the devices continue to exhibit basic function even when a user opts to disconnect them from the rest of the Internet—a practice that many organizations, including the BITAG, have recommended that IoT devices adopt.

Another important aspect of traffic control concerns whether IoT devices on the home network can exchange traffic with one another—and, in fact, whether they can even see one another. Mechanisms, such as network segmentation, which creates virtual “slices” of the same physical network, can ensure that devices are isolated from one another.<sup>22</sup> By default, a user may opt to have new or untrusted IoT devices connect to the home network in a dedicated network slice, to prevent the device from interacting with other devices on the same

---

21. *IoT Home Inspector Challenge*, FTC, <https://www.ftc.gov/iot-home-inspector-challenge> [https://perma.cc/B2H5-H6FU] (last visited Oct. 16, 2017).

22. Yiannis Yiakoumis et al., *Slicing Home Networks*, in *PROCEEDINGS OF THE 2ND ACM SIGCOMM WORKSHOP ON HOME NETWORKS 1–6* (2011).

home network; a user could override this default behavior by explicitly enabling communications between pairs or groups of devices, when such communication is appropriate. Existing technologies such as virtual local area networks (VLANs), as well as emerging technologies such as Software Defined Networking (SDN),<sup>23</sup> may serve as useful building blocks for realizing implementations that achieve isolation between individual IoT devices on the home network.

### *C. Improving the Resilience of the Internet Infrastructure*

Although the Mirai botnet attack garnered widespread attention because it originated from a collection of insecure IoT devices, another lesser noted fact was that the attack's effectiveness was amplified because it targeted not the websites themselves, but rather, the infrastructure that many websites rely on to resolve domain names—specifically, the DNS servers that are responsible for managing the domain names for the corresponding websites. In short, the attack was effective because many of the Internet's websites shared a single common dependency in the DNS infrastructure. These types of dependencies are sometimes (though not always) avoidable; in this case, the vulnerability could have been mitigated had the websites that were attacked had diversified the infrastructure responsible for managing their DNS names (in this case, the solution would have been as simple as using two distinct DNS-hosting providers).

Internet researchers and engineers should be looking for ways to automatically identify—and eliminate—the types of shared vulnerabilities that enabled the Mirai botnet attack. From a technical perspective, such infrastructure diversification is simple; often, it amounts to simply using multiple hosting providers to deliver a particular service. Yet, as a practical matter, these simple measures can sometimes prove challenging, for several reasons. First, there are many opportunities for multiple services to share a single point of failure, and operators may reasonably fail to notice all of these dependencies. For similar reasons, automating the detection of these shared dependencies may also prove challenging. Second, there may be concerns about cost, because using multiple service providers to host a service will inevitably raise the cost of delivering that service. A possible remedy to this state of affairs may be to impose stronger regulations or guidelines about the level of redundancy that infrastructure as a service providers design into their systems.

---

23. Nick Feamster et al., *The Road to SDN: An Intellectual History of Programmable Networks*, 44 ACM SIGCOMM COMPUT. COMMUNIC'N REV. 87 (2014).

## CONCLUSION

Improving the Internet security in the face of insecure IoT devices will require a combination of technical and regulatory mechanisms. Engineers and regulators will need to work together to improve security and privacy of the Internet of Things. As networked devices proliferate and continue to become intertwined with all aspects of our lives—from entertainment to transportation to health—the risks to security and privacy are certain to increase. Other studies such as the BITAG report have catalogued a subset of the security and privacy vulnerabilities, and the potential consequences of those vulnerabilities. Suffice it to say, because IoT devices interact with both the network and critical infrastructure that is connected to the network, vulnerabilities pose substantial risks, from unwanted surveillance to denial of service—both to the Internet services that we have come to depend on and to physical infrastructure that depends on the network.

Engineers must continue to advance the state of the art in technologies ranging from lightweight encryption to statistical network anomaly detection to help reduce risk. Emerging technologies, such as SDN, may help facilitate approaches to network segmentation and home network firewalls that can help users quarantine or control insecure IoT devices without impairing the functionality of the device or other connected devices. Such solutions ultimately require the cooperation of device manufacturers, who must design devices to perform basic functions even when connectivity is disrupted. Similarly, engineers must design the network to improve resilience in the face of the increased risk of attack.

Regardless of the technical solutions that we create, deploying such technology in practice will require the appropriate alignment of incentives so that the parties that introduce risks are more aligned with those who bear the costs of the resulting attacks. Engineers can significantly lower the barrier for device manufacturers by developing certification standards and implementing these standards in easy-to-use, open-source software libraries. Ultimately, however, consumers make decisions about devices to purchase and connect to their networks, and may thus need to be held accountable either directly or indirectly.

