A Vision for Secure IoT

Executive Summary

The rapid proliferation of Internet-connected devices ("Internet of Things" or "IoT") has the potential to transform and enrich our lives and to drive significant productivity gains in the broader economy. However, the lack of sufficient security in these newly connected devices creates meaningful risk to consumers and to the basic functionality of the Internet. Criminals exploit insecure connected devices to create botnets that launch Distributed Denial of Service ("DDoS") attacks against the Internet infrastructure and online services. As seen this past fall, the Mirai botnet used compromised IP cameras and video recorders to launch the DDoS attack that crippled Dyn, a DNS provider, and impaired Internet access to many popular websites for millions of users.

While Mirai perpetrated one of the most impactful recent DDoS attacks, this will surely not be the last event. Such attacks were once the purview of sophisticated hackers. Now, ready-made DDoS "services" are offered to anyone willing to enter the "Dark Web" with Bitcoin to spend. Ever-increasing broadband capacity, a boon to consumers and the economy, also enables increased volumetric attacks. And, most importantly, the number of attack vectors are growing substantially as IoT devices proliferate – with a doubling or more between 2016 and 2020.

IoT therefore represents the next major axis of growth for the Internet. But, without a significant change in how the IoT industry approaches security, this explosion of devices increases the risk to consumers and the Internet. To reduce these risks, the IoT industry and the broader Internet ecosystem must work together to mitigate the risks of insecure devices and ensure future devices are more secure by developing and adopting robust security standards for IoT devices. Industry-led standards represent the most promising approach to broadly increase IoT security. Given the global and constantly evolving nature of threats, industry must utilize its expertise and reach to develop, adopt, and enforce fundamental IoT security measures.

This paper details the technical areas that must be addressed in an IoT security standard. The paper proceeds in three sections: The first provides an overview of the risks posed by insecure IoT to consumers and the Internet. The second highlights the shared responsibility of the entire Internet ecosystem in addressing these risks through mitigation and prevention. The third section details the technical goals of an industry-led, standards-based approach as well as the governance goals of the development organization.[1] To achieve the needed level of security, an IoT security standard must address: (i) device identity; (ii) authentication, authorization, and accountability (onboarding); (iii) confidentiality; (iv) integrity; (v) availability; (vi) lifecycle management; and (vii) future (upgradable) security. A robust technical standard is necessary but not sufficient. To establish value and credibility in the marketplace, the development organization must be open and balanced, ensure due process and consensus, drive wide-spread adoption of the standard, address the intellectual property rights of participants, and ensure conformity through strong certification testing and enforcement of the standard.

An industry-led, standards-based approach provides the most viable path to meaningfully increasing the security of IoT devices, given the cross-border, global nature of the challenge and the rapidly evolving nature of the technology and associated threats.

Need for Action – A Risk to Consumers and the Internet

Insecure IoT devices pose a risk to both consumers and the basic functionality of the Internet. These risks continue to increase based on organic technology and market trends that are making IoT security a growing problem. Most prominently, connected devices are proliferating in the home as broadband capacity grows. As illustrated in Figure 1, leading industry forecasts of IoT growth demonstrate a consensus that the number of connected devices is poised to grow rapidly, with a doubling or more between 2016 and 2020.

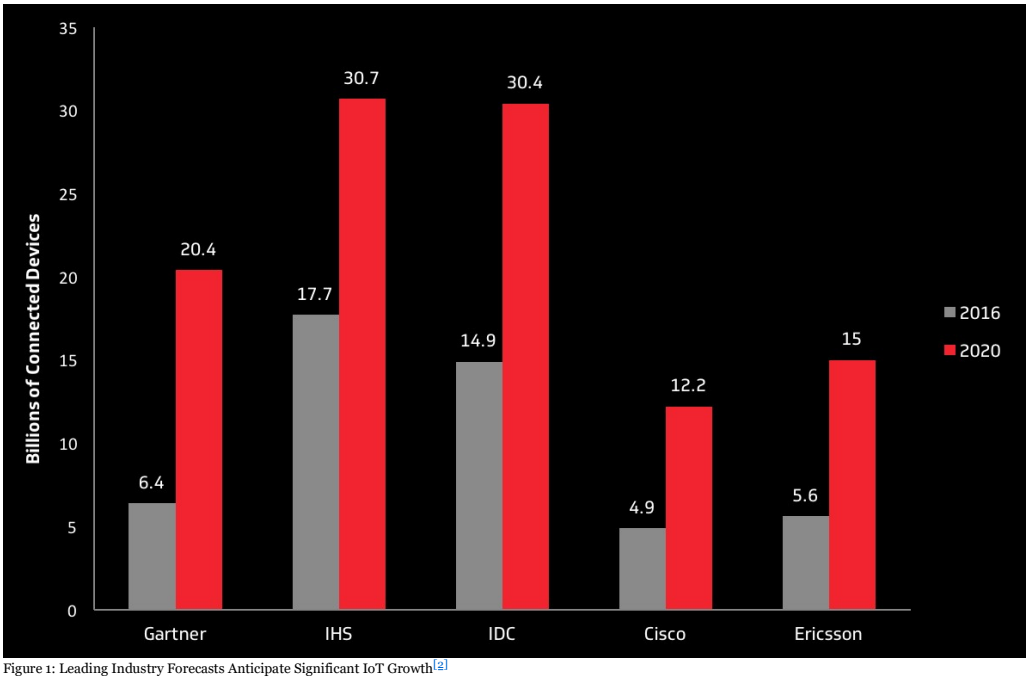

Figure 1: Leading Industry Forecasts Anticipate Significant IoT Growth[2]

Insecure IoT devices serve as the latest building blocks for botnets, which in turn perform DDoS attacks, steal personal and sensitive data, send spam, and more generally, provide the attacker access to the compromised devices and their connections.[3] As depicted in Figure 2, the proliferation of insecure IoT devices enables the emergence of botnets that perpetrate DDoS attacks increasing in both frequency and scale.[4] Moreover, the ready availability and rental of insecure IoT devices for botnet exploit gives rise to "DDoS as a Service," which lowers the barrier – both monetarily and in sophistication – for would-be attackers to deploy such attacks, further driving the increase in DDoS attacks. Cisco explains that the "[a]verage DDoS attack size is increasing steadily and approaching 1.2 Gbps" and "[g]lobally the number of DDoS attacks greater than 1 Gbps grew 172 percent in 2016 and will increase 2.5-fold to 3.1 million by 2021."[5]

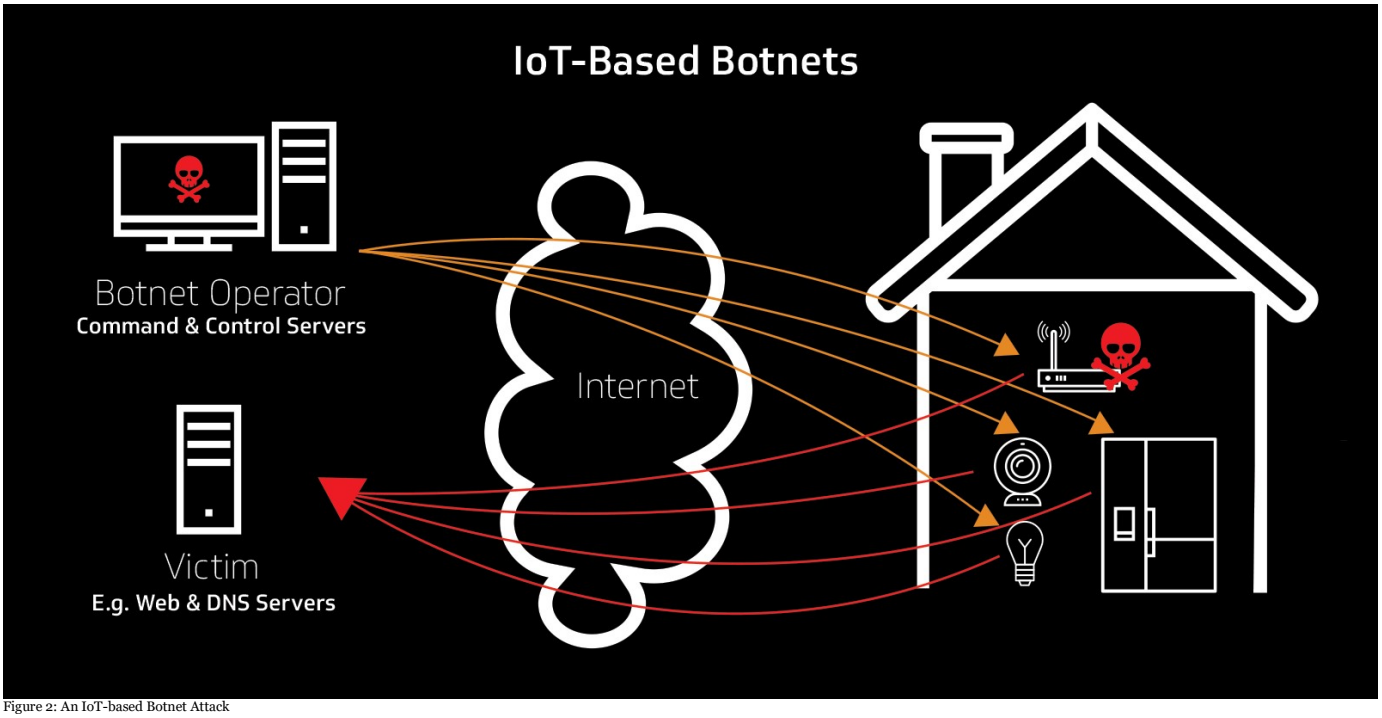

Figure 2: An IoT-based Botnet Attack

Risk to Consumers

Connected devices bridge the physical and cyber worlds, enabling cyberattacks to now cause physical harm in addition to the more traditional privacy and data risks.[6] The physical risks are not hypothetical. In the smart home context, for example, an insecure connected thermostat can expose the consumer to the risk of frozen pipes or other physical damage to his or her home.[7] Researchers have demonstrated the ability to easily unlock a connected door-lock from nearly a half-mile away, enabling a would-be burglar to then walk up to an open door and enter without suspicion or delay.[8] Insecure IoT also increases risk to consumer privacy and the potential of personal and sensitive data theft. This was keenly highlighted in recent allegations that intelligence agencies exploited Internet-connected Samsung TVs to surreptitiously capture audio and possibly video of anything in the room.[9]

The risk to consumers also extends beyond retail devices. For example, researchers have demonstrated the ability to compromise implanted cardiac devices to "deplete the battery or administer incorrect pacing or shocks."[10] Similarly, researchers found a security vulnerability in an insulin pump through which a hacker could cause the pump to deliver a fatal overdose.[11] In the hospitality context, hotel guests were locked out of their rooms when a hacker compromised the electronic key system and demanded payment before returning control back to the hotel.[12]

Risk to the Internet

DDoS attacks, particularly as they grow in both frequency and intensity, raise the real risk of widespread Internet outages. In the Fall of 2016, the Mirai botnet was used to launch a massive DDoS attack against Dyn, a DNS provider.[13] During that attack, many consumers were unable to reach popular sites such as Twitter, Spotify, and Airbnb, causing substantial harm to those businesses.[14] Similar attacks and outages have occurred around the world and the continued proliferation of insecure IoT devices provides potential attackers with the ready fuel to launch the next attack.[15]

The risk to the Internet is not limited to specific geographic regions. Compromised devices from all regions of the world participate in DDoS attacks and other malicious activities. Often the target of an attack is located in a different region than the compromised devices participating in the attack.[16] Improved device security must therefore be addressed from a global perspective and not just country by country.

The real possibility of increasing disruptions to core network functionality and major online service has the potential to undermine public confidence and the increasing role of the Internet in the global economy and society, more generally. This dynamic jeopardizes the benefits of a connected society – civic engagement, digital commerce, and productivity are all put at risk. Fundamental security must be a shared responsibility to ensure continued growth of the Internet's broad-reaching benefits.

Addressing Current and Emerging Threats – A Shared Responsibility

A combination of mitigation and prevention is necessary to more fully address the current and emerging threats posed by insecure IoT. The cable industry recognizes that addressing botnets and other security risks is a shared responsibility across the entire Internet ecosystem. To this end, cable operators have invested substantially in developing and deploying measures to reduce the risks associated with insecure IoT, including DDoS and other botnet attacks, with a primary focus on protecting networks to ensure the availability of broadband service.

Mitigation

Cable operators have developed and continue to improve measures that seek to mitigate DDoS and other attacks against their networks and their customers. These efforts include both individual and collaborative measures, as summarized below.

- **Detection and Identification Systems.** Cable operators have widely deployed and continue to improve systems that are designed to detect compromised customer-owned devices controlled by botnets. These systems rely on (i) high-quality, third-party data feeds that identify sources of malicious traffic on the operator's network, (ii) DNS based anomaly detection systems, (iii) NetFlow detection systems that seek to identify devices communicating with known command and control servers, and (iv) email metadata to identify compromised customer devices originating SPAM. As described below, a DDoS attack source information sharing pilot is also underway to potentially identify sources of DDoS attacks. These information sources are aggregated to confirm whether one or more of a subscriber's connected devices has been infected by a bot.
- **Customer Notification and Remediation Programs.** Once a compromised device is detected and identified, the cable operator will generally seek to notify the affected subscriber and to the extent possible, assist the customer in remediating the compromised devices. Cable operators use a wide array of mechanisms to notify customer of a compromised device, including (i) the use of a captive portal – redirecting the user to a walled garden to provide notification, (ii) in-browser – displaying a browser pop-up box notification, (iii) cable operator website – notification provided when the customer visits the cable operator's website, (iv) email notification, (v) SMS text notification, (vi) app notification, (v) browser toolbar notification, (vi) TV notification, (vii) postal mail, and (viii) telephone notification. To increase notification effectiveness, cable operators are beginning to deploy platforms that allow the customer to select and manage their preferred notification methods. As the number of connected devices in the home increases, identification, notification, and remediation of compromised devices becomes more challenging and cable operators are investigating new approaches and technologies to address these challenges.
- **DDoS Monitoring and Mitigation Systems.** Many cable operators have deployed DDoS monitoring and mitigation systems to ensure the continued availability of their broadband Internet access services during an attack.[17] A DDoS attack seeks to make a device, service, or network resource unavailable to its intended users by flooding the target with superfluous network traffic in an attempt to overload systems and prevent legitimate traffic from getting through to the target of the attack. A significant DDoS attack will typically originate from many thousands or hundreds of thousands of compromised devices. Both the frequency and magnitude of DDoS attacks continue to grow, fueled in large part by the proliferation of insecure IoT. To protect their networks, cable operators typically employ network DDoS mitigation techniques involving specialized equipment that learns and identifies normal Internet traffic sources, destinations and volumes. When Internet traffic anomalies are detected, the abnormal traffic can be separated from the normal traffic so that the DDoS attack does not negatively impact Internet access broadly. However, since attackers are continuously evolving their techniques, these systems cannot provide complete protection.[18]
- **Prevention of IP Address Spoofing.** Source Address Validation (SAV) is a recommended best practice for all ISPs, hosting providers, cloud providers and others to prevent reflective DDoS attacks.[19] SAV with spoofed packet dropping is supported in Cable Modem Termination Systems (CMTS) equipment deployed in cable access networks globally. This feature became available in the Data Over Cable Service Interface Specification (DOCSIS) release 3.0, first issued in 2006, as a mandatory requirement.[20] Moreover, the DOCSIS specification requires that SAV be turned on by default for DOCSIS 3.0 and 3.1 compliant CMTS devices.[21]
- **DDoS Information Sharing Pilot.** CableLabs is working with a number of cable operators and other network operators to develop a DDoS Information Sharing Pilot under the DDoS Special Interest Group (SIG) of the Malware Message Mobile Anti-Abuse Working Group (M3AAWG).[22] This group is developing an API, data store and open source reference implementations for ISPs and other network operators to share information on the sources of DDoS attack traffic. The purpose of this effort is to provide ISPs with actionable intelligence to remediate compromised devices on their networks, rather than mitigating attacks in real time. To generate the actionable intelligence, each participant shares the source IP addresses for the inbound IP flows that their DDoS detection systems identify in an anonymous fashion with the operator of the network on which the DDoS attack originated, as illustrated in Figure 3. To protect privacy, only non-personally identifiable information, such as the source IP address and volume of the attack, are shared.

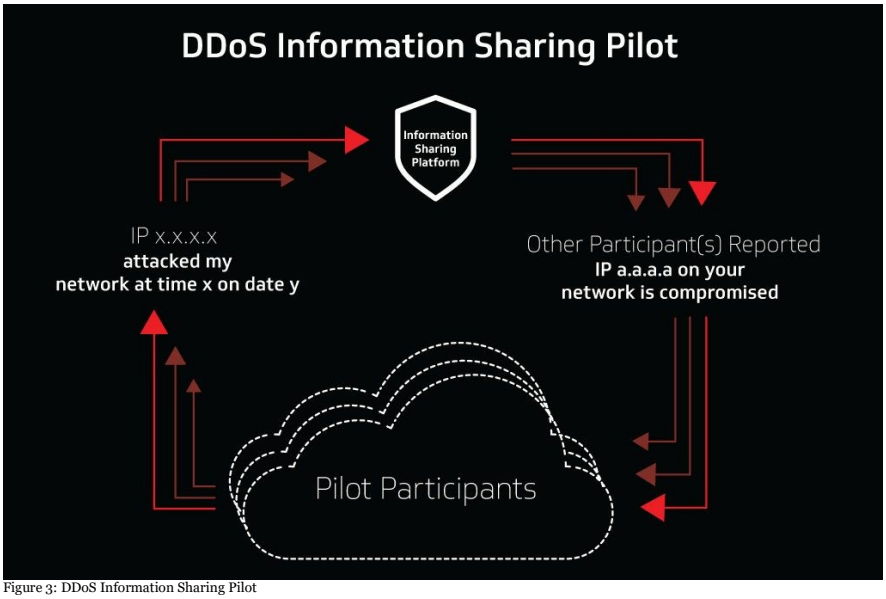

Figure 3: DDoS Information Sharing Pilot

Prevention

An ounce of prevention is worth a pound of cure.
- Ben Franklin (1735)

Although ISPs, including cable operators, have been working on mitigating the effects of compromised and insecure devices for more than 15 years, these efforts ultimately only address the symptoms and not the root cause of the problem. More critically, because of the global nature of the Internet and the ability of a compromised device to target a victim in any jurisdiction, any given ISP has a limited ability to affect the problem beyond addressing the symptoms.[23] Furthermore, the challenge of this task has already begun to outpace current and anticipated mitigation techniques as the number of connected devices is expected to double by 2020.[24] Unfortunately, IoT providers have not generally incorporated reasonable security measures or committed to maintaining the security of their IoT devices. In contrast, the major computer operating system manufacturers in recent years have significantly increased their efforts (including timely security patches) to ensure general-purpose computing devices (e.g., desktops, laptops, tablets, and smartphones) are reasonably secure from malicious software. This lack of security in the growing population of IoT devices is increasingly affecting the overall security of the Internet.

To more fully address the risks posed by insecure IoT devices, industry must drive increased security into future IoT devices. Preventing compromised devices must be a substantial part of the industry's shared responsibility in addressing the risks posed by insecure IoT to consumers and the Internet. To minimize these risks, device security must be improved in the areas detailed below, through an industry-led, standards-based approach.

A Vision for Secure IoT: Increasing Security through an Industry-Led, Standards-Based Approach

To help stem the tide of insecure IoT devices and help address the above risks, industry must work to develop and adopt the necessary standards to ensure connected devices have incorporated sufficient security. For an industry-led, standards-based approach to be credible and succeed, it must be (i) robust, (ii) broadly adopted, and (iii) have a strong certification testing and enforcement mechanism.

A comprehensive, industry-led approach to IoT security must account for both the technical security features of the device and ecosystem as well as the governance of the organization developing, certifying, and enforcing the standard. The technical security goals detailed below are a base level of security, envisioned for every device connected to the Internet – either directly or indirectly. We focus on IoT devices sold to and intended for the retail, consumer market. Additional security features may be necessary for devices used in more sensitive applications, for instance, in healthcare, industrial, or transportation. Our technical goals are drawn from the cable industry's experience and the work of industry stakeholders, including Open Web Application Security Project (OWASP), Broadband Internet Technical Advisory Group (BITAG), Wi-Fi Alliance, GSMA, AT&T, Microsoft, and Cisco, as well as the work of government agencies, including the Federal Trade Commission (FTC), National Institute of Standards and Technology (NIST), and the Department of Homeland Security (DHS).

The substance of the security standard is critical but how the standard is developed, maintained, and enforced is also critical. The development organization must have the widely-accepted governance attributes set forth below to ensure credibility and value to buyers and the Internet ecosystem, more broadly. Our governance goals are drawn from the work of both public and private organizations, including the U.S. Office of Management and Budget (OMB), British Standards Institution (BSI), International Organization for Standardization (ISO), International Electrotechnical Commission (IEC), the World Trade Organization (WTO), and the Standards Council of Canada.

Technical Security Goals

For any industry-led, standards-based security approach to be effective and successful, it must be robust and materially improve IoT security. The lapses in IoT security to date are well documented.[25] To meaningfully and comprehensively improve IoT security, a security standard must address the following areas. A summary of the technical security goals is provided in Appendix 1.

To support strong security, the device identifier must be immutable, attestable, and unique.

Device Identity

Device identity is the foundational building block for IoT security – it is the necessary prerequisite for many other crucial security elements. But, not all device identifiers are equal. To support strong security, the device identifier must be immutable, attestable, and unique. Today, IoT devices typically do not use identifiers that are both unique and immutable and the device identifiers are almost never attestable. Attestability enables the device identity to be cryptographically verified, dramatically reducing the risk that the device is being impersonated (or "spoofed"). As discussed in more detail below, a strong device identifier (i.e., immutable, attestable, and unique) is the necessary foundation to building secure IoT devices.

The technology exists to provide strong device identity that can be scaled to meet the demands and device volumes anticipated with IoT. The cable industry has a long incorporated strong device identity in its cable modems, set-top boxes, and other devices directly connected to the cable network, using a public key infrastructure (PKI). The cable industry has issued more than 500 million device and software code validation certificates. This same technology and infrastructure can readily scale to meet the device volumes anticipated with IoT.[26]

Critically, certificate management is the mechanism that enables attestation, which not only enables verification, but also revocation – allowing the certificate manager to deprecate or completely revoke a certificate and communicate that revocation in response to future inquiries. In turn, certificate management enables enforcement of security standards among and between devices, as discussed in more detail below. The certificate manager can serve as the authoritative source as to whether a device has passed the certification test associated with the security standard. For example, to protect itself, a connected healthcare device might only communicate with other devices that have been certified to meet a higher level of security within the standard. The healthcare device can query with the certificate manager (or certification authority), using the digital certificate of the other device, to verify whether that device actually complies with the necessary security standards.

Certificate management and the ability to revoke certificates also provides an automated mechanism to ensure ongoing compliance with a security standard or to communicate compromises or known vulnerabilities for devices. For example, if a device initially passes certification testing, but a critical security vulnerability is discovered or the manufacture then makes changes (e.g., adds additional features) that cause the device to no longer comply with the security standard, the certificate manager can revoke those certificates until the manufacturer addresses the issue. More generally, with a certificate manager, anyone (the ISP, another device, or smart home hub) can query whether a device is and remains compliant with the security standard and if not, deprecate that device's privileges.

Authentication, Authorization, and Accountability - Onboarding

Secure authentication, authorization, and accountability minimize the potential for compromising a device or other devices in the local IoT ecosystem during the onboarding process. "Onboarding" is the process by which a new device is connected and added to the network and the local IoT ecosystem. Onboarding includes the processes for authentication, authorization, and accountability (AAA) of that new device. Authentication is the process by which the device identity is verified and confirmed. Authorization determines what network resources the device will have access to. And, accountability is the process that tracks what the device does.[27]

The use of secure digital certificates enables a more straightforward and secure onboarding process for devices. The exchange of secure digital certificates between the new device and an IoT hub or other devices on the network allow the devices to determine the level of trust between devices and what information to share. This is achieved by the devices exchanging certificates and each device confirming with the certificate authority on the status of the other devices. Use of certificate-based onboarding not only increase security, but it also can enable an easier, more customer friendly onboarding process (e.g., no confusing PIN to enter). For example, the Wi-Fi Alliance has incorporated a centrally managed certificate-based approach for onboarding Wi-Fi connected devices in their Passpoint specifications.[28] This digital-certificate approach eliminates the need to remember and enter passwords as part of the authentication process.

Insecure onboarding can introduce vulnerabilities to the new device, other devices on the network, and the network itself. These vulnerabilities can be exploited through "man-in-the-middle" and reprovisioning attacks as well as device spoofing and snooping. The use of secure digital certificates minimizes the potential for onboarding vulnerabilities.

Confidentiality

Strong confidentially protections ensure sensitive information remains private and inaccessible to unauthorized parties. Ensuring the confidentiality of sensitive information goes beyond just encryption. IoT devices should protect sensitive data at rest, in use, and in transit and limit the information disclosed in response to anonymous or untrusted requests. The IoT device manufacturer must first identify the sensitive information a device handles. This may include personally identifiable information (PII), protected health information (PHI), credentials, and private keys, to name just a few categories.

Protecting Sensitive Information (In the Device): First, an IoT device should encrypt locally stored sensitive information, using standard, well-established encryption techniques.[29] Second, the device should protect sensitive information (e.g., private keys) while in use in the device. For instance, private keys should reside in dedicated hardware and not be transmitted in the clear across the bus.[30]

Protecting Sensitive Information (In Transit): An IoT device should use mutual end-point authentication and application-level encryption (end-to-end) for all communications that include sensitive data. Each device in an IoT ecosystem should authenticate all other devices that participate in that ecosystem. Once authenticated, each device would encrypt and sign messages sent to other devices in the network. Each device that receives a message can then cryptographically validate the data prior to acting on it. However, mutual authentication does not address the security of the sensitive information as it is passed beyond the authenticated devices and local IoT ecosystem. Therefore, IoT providers should also incorporate application-level encryption (end-to-end) to ensure sensitive information is securely passed between the IoT device and the servers owned by the IoT provider.[31]

Limiting Responses to Untrusted Requests: An IoT device should limit the information provided in response to requests by untrusted (e.g., anonymous) devices. Prior to onboarding, an IoT devices should use a temporary, ephemeral random identifier in response to new onboarding requests. After onboarding, the device must provide its immutable, attestable, and unique identifier, as discussed above. Providing a device's unique and immutable identifier to any and all requests creates security risk. For example, a mobile device, such as a Bluetooth fitness band, that broadcasts its unique and immutable identifier whenever requested, enables easy tracking of the device

and its owner without the owner's knowledge or consent.[32]

Integrity

The data created or received by a device must be trustworthy, and protected from unauthorized modification. This requires that the device identity, execution environment, configuration, and communications are secured using well-established methods.

> The unwitting foot soldiers of Mirai, insecure webcams and DVRs, were largely compromised by scanning for open Telnet ports on these devices and exploiting those open ports through the use of hard-coded credentials.

To further ensure device integrity, each device should be "hardened" to minimize the attack surface by closing unnecessary ports, disabling unnecessary services, and using a secure bootloader with configuration validation.[33] The Mirai botnet provides an illustrative case study. The unwitting foot soldiers of Mirai, insecure webcams and DVRs, were largely compromised by scanning for open Telnet ports on these devices and exploiting those open ports through the use of hard-coded credentials.[34] In addition to not incorporating hard-coded credentials, the manufacturer should consider disabling unnecessary ports and services to avoid increasing the risk of vulnerability and exploit.[35]

Manufacturers should also employ non-repudiation methods for critical communications to ensure the integrity, origin, and/or delivery of the data as well as a trusted audit trail to establish that data was sent and received in an unmodified manner.[36] Non-repudiation methods are particularly important when a direct financial consequence is connected to the data. For instance, in the smart electrical meter context, non-repudiation methods are employed to prevent both an end-customer denying sending energy consumption data and the electric utility denying sending real-time pricing data.[37]

Availability

A secure IoT device is available when it is needed for its legitimate use and unavailable when it is not. IoT devices should be designed to function in a predictable and expected manner, if and when there is a loss of broadband connectivity or a loss of communications with any associated cloud service. Conversely, devices should use restrictive, rather than permissive, default network traffic policies to limit communications to expected norms, guarding against both unintended as well as malicious denial of service attacks that can disrupt the availability of the device or other devices on the network.

Specifically, a device should limit the information provided in response to discovery and other requests from untrusted sources. For example, a device should provide limited, if any, information in response to reflection and introspection requests from an untrusted source. Such requests, particularly in a repeated fashion, can be in effect a denial of service attack against the device that is the target of these requests. Limiting responses to untrusted requests will also help prevent a rogue or compromised device from gaining information about other devices on the network as part of a "stepping stone" attack.

Manufacturers should also ensure connected devices will continue to function, in an expected manner, in the event of short-term disruptions of connectivity or longer-term outages. These disruptions or outages may occur in the local-area or access (broadband) networks or with the cloud service. Disruptions and outages will occur; the question is how will the connected device react. For instance, consumers should expect a connected thermostat to continue to control the heating and cooling in their homes even if connectivity to the cloud service is lost for an extended period of time.[38] Similarly, the manufacturer of a home alarm system should make clear to the consumer how the system will respond (e.g., an alarm will sound) to a loss of connectivity to either the cloud service or to sensors or other peripherals in the system.[39]

Lifecycle Management

IoT security requires vigilance throughout the life of the device – vulnerabilities will be discovered and new threats will emerge after the consumer purchases the device. IoT providers must make lifecycle management a central consideration in the design of every connected device and clearly disclose the key considerations to consumers prior to sale.[40] Specifically, IoT providers must, with limited exception for ephemeral devices, provide secure, automated, software updates during the disclosed security support period. In addition, IoT providers must publicly disclose vulnerability remedies and changes to functionality at end-of-life (EOL)/end-of-support (EOS).

Software Updates: IoT providers must provide secure, automated software updates throughout a clearly defined and disclosed security support period.[41] By default, the software update mechanism should not require or rely on any consumer action.[42] IoT providers incorporating a secure, automated software update mechanism into their devices recognize the reality that vulnerabilities are discovered in devices after they are deployed and that software updates can mitigate the risks associated with these vulnerabilities.[43]

To ensure secure software updates, the IoT provider must use cryptographic checks to ensure the origin and integrity of the software update. This is another area where an IoT provider can leverage PKI. For example, cable operators use their PKI to securely provide software updates to cable modems and other cable devices. The digital certificates ensure that only software updates from either the manufacturer or from the cable operator can be downloaded into the cable device and the certificate is also used to encrypt the update to ensure its integrity. The use of digital certificates in this manner has minimized the risk that cable devices will be infected with malware or other malicious code as part of the software update process.[44]

We recognize that not all IoT devices require a mechanism for software updates. In particular, to address vulnerabilities in devices that are very inexpensive and/or have a very limited life, replacement may be a more economical alternative to providing software updates – for instance, in disposable wireless medical bandages.[45] However, for such devices, the IoT provider should have a mechanism to identify vulnerable devices, disable vulnerable devices, and communicate the need for replacement of vulnerable devices to end-users.[46]

Vulnerability Management: An IoT provider should have a well-defined procedure for receiving reports of security issues for their devices. The procedure should include status reporting and a timeline to address the problem that is provided to the individual or entity that submitted the security vulnerability. At a minimum, the IoT provider should publicly and prominently disclose an email address, a telephone number, and a website where security issues can be submitted to the company. Once there is a remedy to the vulnerability, the IoT provider should have a mechanism to publicly disclose the vulnerability and associated remedy.

End-of-Life (EOL) / End-of-Support (EOS) Functionality: To protect end-users and third-parties, IoT providers should consider limiting device functionality after the security support period ends. Prior to sale, IoT providers should clearly disclose whether and to what extent device functionality will be limited due to an increased risk of vulnerability after the security support period ends.[47] To set consumer expectations, the disclosure should describe exactly what, if any, functionality will be limited at the end of the support period – whether only the "smart" functions and features (e.g., connectivity and control remotely through an app) will become inoperable, or whether core device functionality will be lost as well.[48]

Future (Upgradable) Security

IoT providers should consider and design into their products the ability to have strong security controls including secure cryptographic algorithms/cipher suites for the entire intended and expected life of the device. A device with a short lifespan (e.g., less than one year) may not require the capability to upgrade. In comparison, providers of connected, durable home appliances (e.g., expected service life of 10 or more years) should consider how the security controls will need to evolve over the life of the device. Manufacturers should evaluate both software and hardware requirements to ensure the upgradability of security controls based on the intended and expected life of the device and the currently employed security controls. For example, home appliances deployed today using recommended cryptographic algorithms will most likely need to be upgraded to stronger cryptographic algorithms and longer key lengths during the expected life of the appliance. Manufacturers should ensure adequate hardware capabilities are included in the appliance for these anticipated upgrades.

Governance Goals

> The substance of the security standard is critical, and how the standard is developed is also important to how that standard is received and adopted by end-users and IoT providers alike.

In addition to ensuring a technically robust security standard, the development organization must employ the widely-accepted governance attributes in developing the standard to ensure credibility and value in the marketplace for IoT devices and in the broader Internet ecosystem. The substance of the security standard is critical, and how the standard is developed is also important to how that standard is received and adopted by end-users and IoT providers alike. For ease of reference, a summary of the governance goals detailed below is provided in Appendix 2 to this paper.

Openness

Industry consensus standards are best developed through transparent processes, open to interested parties to meaningfully participate on a non-discriminatory basis.[49] This should include openness with respect to participation at the policy development level and at every stage of the technical standard development. The development organization should also seek to ensure participation of interested parties with limited resources – for example, civil society.[50]

Balance

To truly capture the broad base of industry expertise, different perspectives must be represented in the standards development process. The development organization as well as the development process for the standard must have meaningful involvement from a broad range of parties, with no single interest dominating the decision-making. All interested parties should be provided with meaningful opportunities to contribute to the elaboration of the standard to ensure balanced representation of interested parties including, but not limited to, the categories of producers, suppliers, buyers, and consumers.[51]

Due Process – Notice, Transparency, Appeals

Participation must also be equitable for those involved. The development organization must adhere to the basic notions of due process, including notice, transparency, and appeal procedures. Notice and transparency require that interested parties have access to written policies and procedures, adequate notice of meetings and standards development, sufficient time to review drafts and prepare views and objections, and access to views and objections of other participants.[52] The development organization must also have a fair and impartial process for resolving conflicting views and for resolving substantive and procedural appeals.[53]

Consensus

Decisions must also be reached in a fair manner. The development organization must use a consensus decision-making process that involves interested stakeholders. "Consensus" is defined as "general agreement, but not necessarily unanimity."[54] The organization must consider comments and objections using "fair, impartial, open, and transparent processes."[55]

Adoption

Wide adoption is critical for a security standard to meaningfully improve the security of IoT. The development organization should not only seek to reduce the barriers (e.g., cost and labor) to adoption but also actively encourage adoption by IoT providers. For example, the development organization should actively engage in end-user education to ensure consumer awareness of the standard and the related certification mark(s) and the value of the increased security that accompanies incorporation of the standard. Creating consumer demand will accelerate adoption by IoT providers.[56]

Intellectual Property Rights (IPR) Policy

Intellectual property rights must be clearly addressed upfront both to ensure broad participation in the standards development process and to reduce the barriers to adoption of the standard. The development organization must have an intellectual property rights (IPR) policy that requires participants to (i) disclose any necessary patents and (ii) license those patents to implementers of the standard on non-discriminatory and royalty-free or reasonable royalty terms (and to bind subsequent owners of standards essential patents to the same terms).[57] The IPR policy should be easily accessible, set out clear rules governing the disclosure and licensing of the relevant intellectual property, and take into account the interests of all stakeholders, including the IPR holders and those seeking to implement the standard.[58]

Conformity Assessment -- Certification Testing and Enforcement

In order for consumers to have appropriate security assurances, there must be a mechanism to validate device compliance with security standards. Rigorous conformity assessment through certification testing and enforcement are critical to any security standard's credibility and value in the marketplace and an industry-led standard's ability to enable competition and market forces to help drive improved security in connected devices. The development organization must ensure that the products claiming to incorporate the security standard actually conform to that standard as a condition of using the associated certification mark. To ensure conformity, the development organization must adopt a strong and transparent certification testing and enforcement program, including advanced and post-hoc compliance testing and a mechanism for withdrawing product certification as appropriate.[59]

> Without rigorous conformity assessment, the certification mark provides a hollow promise and little, if any, meaningful information to buyers, fundamentally undercutting the value of the standard and the associated certification mark.

Backed by rigorous conformity assessment, the certification mark provides potential buyers with the information and confidence that the product meets or exceeds the minimum-security features and controls set forth in the

standard. Without rigorous conformity assessment, the certification mark provides a hollow promise and little, if any, meaningful information to buyers, fundamentally undercutting the value of the standard and the associated certification mark. To be clear, we do not believe that governments should mandate a security standard or certification testing regime in the context of consumer connected devices.[60]

Conformity assessment is another area where strong device identity can be leveraged to ensure compliance with the standard. Through attestation and revocation, the network and other devices can use a device's identity (e.g., digital certificate) to query whether that device has been tested and certified compliant with the security standard and whether the device remains compliant or had its certification revoked. In onboarding a new device, the local network or other devices can use the compliance information provided by the certificate manager to determine the level of trust to afford to that new device. This use of strong digital certificates with attestation eliminates the ability of IoT providers to make false claims around standards compliance and minimizes the burden on consumers in onboarding and policing device compliance.

Conclusion

The rapid proliferation of insecure connected devices is increasing the risk to consumers and to the basic functionality of the Internet. To reduce these risks, the IoT industry and the broader Internet ecosystem must work together to mitigate the risks of insecure devices and ensure future devices are more secure by developing and adopting robust industry-led security standards for IoT devices. As detailed in this paper, an industry-led, standards-based approach must comprehensively address the technical areas of security and ensure the development organization is open and balanced, ensures due process and consensus, drives wide-spread adoption of the standard, addresses the intellectual property rights of participants, and ensures conformity to the standard through strong certification testing and enforcement.

Appendix 1: Summary of Technical Security Goals for IoT

Category	Description	Goal for Every IoT Device
DEVICE IDENTITY	Require an attestable, immutable, and unique identifier for each device.	Use a secure certificate-based approach (PKI) to provide an attestable, immutable, and unique identifier for each device; certificate issuance; lifecycle management; and revocation.
AAA/ONBOARDING	Require the use of strong Authentication, Authorization, and Accountability (AAA) methods for provisioning and management of devices.	Use strong device identity to enforce strong authentication (identifying the user or device), no shared default credentials; use authorization enforced through established mechanisms to provide access to network and device resources; incorporate accountability mechanisms that associate device actions with authentication and authorization.
CONFIDENTIALITY	Protect sensitive data from access by unauthorized individuals, entities, devices, or processes.	Identify sensitive information (PHI, PII, credentials, etc.) and protect that data appropriately. Use of encryption for locally stored sensitive information. Provide protection of sensitive information (e.g., private keys) while in use. Use mutual end-point authentication and application-level encryption (end-to-end) for sensitive data in transit. Provide an ephemeral identifier for anonymous discovery requests and limit information available to anonymous introspection and reflection requests.
INTEGRITY	Assure the device is trustworthy and the processes, data, and communications associated with the device are accurate.	Confirm that the device identity, execution environment, configuration, and communications are authorized and appropriate, using well-established methods. Harden the device to minimize the attack surface by closing unnecessary ports, disabling unnecessary services, and using a secure bootloader with configuration validation. Consider use of non-repudiation methods for critical communications.
AVAILABLILITY	Safeguard devices and associated communications for proper functioning.	Plan for appropriate device behavior in the event of network or radio communication interruptions or outages. Use restrictive, rather than permissive, default network traffic policies to limit communications to expected norms, guarding against both unintended as well as malicious denial of service attacks.
LIFECYCLE MANAGEMENT	Support sufficient secure operation, update, and communications throughout the life of the device.	Provide for secure, automated, update mechanisms during the disclosed support period and publicly disclose vulnerability remedies and EOL/EOS functionality changes.
FUTURE (UPGRADEABLE) SECURITY	Plan for security improvements required to support equivalent device and network security in concert with Lifecycle Management.	Include support for longer key lengths, stronger cryptographic algorithms/cipher suites, and hardware based security over the supported life of the device.

Appendix 2: Governance Goals for Development Organization

Category	Description	Goal for Development Organization
OPENNESS	Criteria to participate.	Open to all interested parties on a non-discriminatory basis, including by those with limited resources (e.g., civil society).
BALANCE	Degree of representation and participation from across the full ecosystem.	Provide balance of interest. Avoid dominance from a single interest category and ensure fair and equitable representation and participation from the full ecosystem of stakeholders.
DUE PROCESS – NOTICE, TRANSPARENCY, APPEALS	Documented policies and procedures for notice, transparency, standards development process, and appeals.	Documented and available policies and procedures. Provide timely and adequate notice of meetings, and standards development activity. Access to views and objections of other participants. Fair and impartial process for resolving conflicts, including procedural appeals. Availability of standards and conformity assessment procedures.
CONSENSUS	Decision process.	Decisions should be reached through a collaborative process incorporating input from all interested parties. Consensus means general agreement, not necessarily unanimity.
ADOPTION	Use and acceptance of the standard and conformity assessment process.	Broad adoption of the standard and conformity assessment process across the ecosystem. Consumer-facing information should convey security reliability.
INTELLECTUAL PROPERTY RIGHTS	RAND or RF patent policy.	Require disclosure of essential patents and adherence to RAND or RF terms of license by participants.
CONFORMITY ASSESSMENT - CERTIFICATION TESTING & ENFORCEMENT	Mechanism to ensure compliance with the standard.	Strong and transparent certification and enforcement program, including advanced and post-hoc compliance testing and a mechanism for withdrawing certification as appropriate.