Comment in response to Dept. of Commerce NTIA "Request for Comments on the Benefits, Challenges, and Potential Roles for the Government in Fostering the Advancement of the Internet of Things"

We appreciate the opportunity to provide comments to the NTIA on the above referenced matter, which encompasses a wide range of stakeholders and touches on many safety, security, and privacy topics. The diversified nature of the ecosystem and technologies would benefit from a multistakeholder process.

This comment is not meant to be comprehensive--it will not iterate the multitudes of benefits that could be realized with the Internet of Things. Here, we want to address the risks and challenges in the safe and secure design of the IoT future, and highlight some opportunities for a regulatory path forward.

1. Are the challenges and opportunities arising from IoT similar to those that governments and societies have previously addressed with existing technologies, or are they different, and if so, how?

The challenges and opportunities arising from the IoT are fundamentally new, owing to the unprecedented combination of ubiquity, diversity, and connectivity among IoT devices, and the ability for many IoT devices to observe and actuate real world events without any explicit human interaction. Despite this, we believe that the challenges are best addressed by focusing on the age-old engineering design principle of "least surprise": The IoT should behave in a manner that is both expected by and clearly communicated to every stakeholder with which it interacts. Realizing this ideal will require a marriage of traditional security- and privacy-enhancing technologies like cryptography and access control with qualitative assessments of diverse stakeholder requirements.

We illustrate the unique challenges raised by IoT by recalling a few recent high-profile vulnerabilities in the IoT: refrigerators that botch certificate verification and can thereby enable attackers to steal your Google password; webcams that ship with unsafe default settings and passwords and can thereby unknowingly broadcast your nursery or bedroom to strangers on the Internet; eldercare door locks (designed to prevent wandering) that default to locked when power is lost and can thereby trap your elderly parents in their home in the event of a fire; and IoT light bulbs that can be made to overheat by a remote attacker and can thereby set your house fire. Unlike laptops, tablets, and smartphones--which are typically owned and operated by a single user, who thinks of them as "computing devices", and is therefore imminently cognizant of the potential for security failures--we believe that it is unrealistic to expect the users of IoT devices to anticipate and mitigate against such risks. Indeed, many stakeholders affected by such devices may never even be aware of the devices existence until a harm occurs.

Other latent vulnerabilities emerge in the context of use. For example, in facilities where individuals are specifically identifiable, changes in health (ranging from prostate cancer to

pregnancy) can be inferred solely from changes in water usage, leading to unanticipated privacy breaches when smart meters are used. In this example, as above, transparency and privacy policy negotiations are key issues as users may be completely unaware that such detailed data are being collected or what sorts of sensitive inferences this data may entail.

The deceptively simple cases mentioned above illustrate the manner in which technical challenges around threat assessment, preference diffusion, networking, and cryptographic enforcement are all intertwined with fundamental physical safety and foundations of personal privacy. These illustrate how the IoT is dissimilar to other technologies. Thus, while we believe that the principle of least surprise will be paramount to addressing the challenges raised by IoT, we stress that an instantiation of this principle for IoT will necessarily look very different than it has for existing technologies.

2. The term ‘‘Internet of Things’’ and related concepts have been defined by multiple organizations, including parts of the U.S. Government such as NIST and the FTC, through policy briefs and reference architectures. What definition(s) should we use in examining the IoT landscape and why?

It is useful to think of IoT technologies as networked technologies which are aware, active, and adaptive.

- *Aware technologies* are capable of sensing their environment; for example, an IoT device may know when a door is opened and closed. Aware technologies do not require explicit user interaction to obtain such data.
- *Active technologies* can respond to events; for example, the above door sensor may be deployed in the home of an older adult experiencing early-stage dementia in order to generate alerts and notify caregivers if the front door is opened late at night (a possible indication that the older adult is wandering from the home). Active technologies do not require specific decisions on the part of user to act.
- *Adaptive technologies* can change with the individual, learning new patterns to generate reminders and alerts and thereby reducing the number of false positives. For example, the above elder may have dinner parties that occasionally extend late into the evening, so the system learns to not generate alerts until after the last visitor has departed for the evening. Adaptive technologies may be a black box, creating situations in which individuals cannot assert specific preferences even if they can articulate them.

All four components--networked, aware, active, and adaptive--are necessary for a definition to be adequate.

6. What technological issues may hinder the development of IoT, if any?

The tendency to seek isolated solutions is inadequate with the IoT. Effective risk mitigation requires a system-wide view of technical issues and points of interoperability, one that encompasses the interactions we illustrated above. In short, a holistic view is needed.

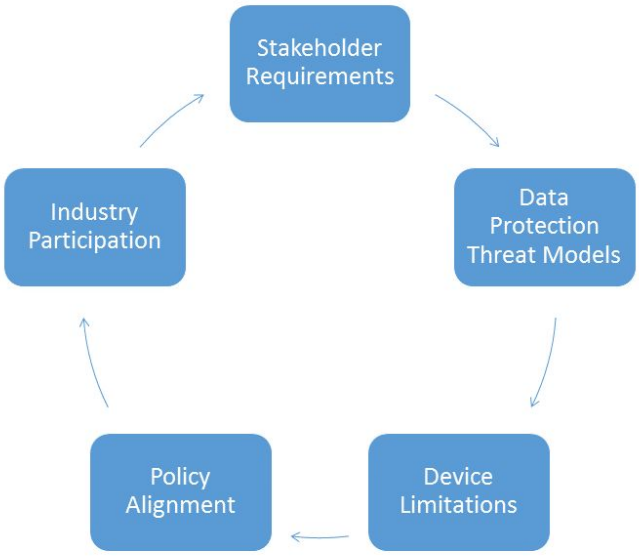

When designing the computing architectures comprising the IoT, taking a system-wide view and identifying the gaps is a foundational challenge. Any effective research agenda must address the interplay between human-computer interaction (HCI), networking and security, data protection and cryptography, and distributed systems. In this, one of the largest gaps is a lack of human-centered research. Government efforts, such as those by DARPA and DHS, have supported great advancements in algorithmic and big data approaches to security, but this has come at the expense of HCI and user studies. An emergent challenge with the IoT is that IoT systems are not merely used, but lived with. This means the platforms come into contact with a diverse collection of stakeholders, which may include vulnerable populations such as children and persons that rely on in-home care technologies. There are concluded initiatives that inform the understanding of these issues in the policy landscape.

Qualitative investigation of stakeholder requirements must inform the threat models that data protection policies and cryptography work to protect against. Without informed threat models, these solutions risk solving the wrong problems. In turn, these functions must be customized to operate with heavily constrained computing resources. Without customized solutions for the devices, data protection solutions may not function within the severe CPU, power, and other constraints typical of IoT devices.

To design for this complex environment and these requirements, a holistic view is necessary in the developing the research that will inform the design of a solution. Without the cryptography, networking, and devices working together, the security and privacy requirements of the stakeholders cannot be enforced. As many risks will emerge only when IoT devices are

deployed in new combinations and use cases, privacy and security protections must be built-in using the aforementioned design principle of least surprise.

16. How should the government address or respond to cybersecurity concerns about IoT?

16(a). What are the cybersecurity concerns raised specifically by IoT? How are they different from other cybersecurity concerns?

Many modern techniques in computer security do not naturally port to many IoT devices, which will have significantly smaller computing capabilities, storage and memory, network connectivity, and user interfaces. This means that many devices will be constructed without the mainstays of modern computer security, and in particular some of the protections against “buffer overflows” and related attacks, which make up the majority of worrisome attack patterns on modern day systems. For example Address Space Layout Randomization (ASLR) will not work in many systems because it requires large address spaces that will not be present. In addition, even within the smaller address space, many devices will not support memory architectures in which memory segments can be marked as non-executing. As well, few modern and emerging cryptographic techniques are computationally efficient enough for computationally-limited IoT devices. For example, it is unlikely that emerging lattice-based cryptographic standards that seek to provide “post-quantum” computing security can be implemented in a light switch.

Entirely new types of threats will emerge in the IoT context. For example, attacks over large numbers of diverse devices may have cumulative effects that individual attacks might not. For example, an attacker who can “brick” an IoT-enabled furnace in the middle of winter might cause inconvenience (and cost, with need to replace the furnace and possibly a stay for the homeowner in a hotel for several nights). An attacker that can brick heating systems for large segments of a community during extreme weather conditions might require a response so large as to impinge the National Guard. Attackers could create conditions where large numbers of people needing to leave their homes at the same time could easily overwhelm local emergency response, support for the disabled, even commercial facilities (such as lodging capacity). Similar concerns have been raised with ability to brick large numbers of smart-grid power meters, traffic managements systems, etc. Indeed, in these latter cases, the risk is amplified by the homogeneity of such systems in a given region, induced by organizational and economic factors (e.g., every house in a neighborhood is likely to have a smart power meter from the same electrical company, installed in the same time frame).

A further concern with IoT devices is that many of them will be sold by retailers and sellers that are not accustomed to the need for constant updating and patching. For example, most retailers might sell a fridge or furnace with a 5 year warranty, but the expectation is that the device will be in use for much longer periods, up to 20 years. If software in old IoT appliances are not patched on a regular basis after the traditional warranty of a device expires, then homes, businesses, and infrastructure will become minefields of embedded insecure devices. Note that we already

seen some similar issues to this problem in other areas of embedded systems, from the United States Navy's need to purchase extended support for Windows XP, to the systematic dependence of ATMs on that same version of Windows, and even the lack of support for individual consumer's modern smartphones by the major providers.

All of the above suggest that new methods and approaches are going to be necessary for securing the IoT landscape. Several of these differences suggest an opening for policy to help direct and coordinate solutions, and consider potential responses to bad outcomes that result from attacks which result in exploitation of the above problems.

16(c). What role or actions should the Department of Commerce and, more generally, the federal government take regarding policies, rules, and/or standards with regards to IoT cybersecurity, if any?

The challenge in realizing this necessary confluence presents three opportunities for the public sector to guide the way forward. The first opportunity is as a convener and supporter of innovation. The multistakeholder process is one such venue to bring elements representing different aspects of the IoT ecosystem into one place.

The second opportunity for the public sector is as a standards setter. Here, the process can begin with minimal safety standards, and move toward addressing best practices in accordance with design principles such as least surprise. There are many existing frameworks to inform such an effort, and in some cases, these minimum standards already exist. The initial motivation for public sector engagement in standards setting is that with the IoT, security problems translate to real-world safety problems. In recent years, regulators such as the Federal Trade Commission (FTC) and the Federal Communications Commission (FCC) have exhibited the ability to address egregious privacy situations that put consumers in risk of real harm. This same regulatory function would enable enforcement of minimum standards, and identify or sanction players who operate in bad faith or push goods of bad quality into the marketplace (such as the recent hoverboard incidents, which elicited safety standards from the U.S. Consumer Product Safety Commission). Manufacturer assertions about quality--including security and privacy--can be enforced through regulation of deceptive or unfair trade practices (Section 5 of the FTC Act), or by various other mechanisms such as the legal requirements to protect confidential customer information (Section 221, 222 of FCC authority).

The third opportunity for the public sector is to address the lemons market in security and privacy. By supporting transparent and comprehensive consumer education, and holding producers accountable for their assertions about quality, consumers can be enabled to make informed decisions, and make security and privacy choices in accordance with their preferences. Consumers should be able to trust manufacturer's assertions with respect to security and privacy. Manufacturers and producers should be confident that there are no cryptographic backdoors in standards intended for IoT use, in no small part due to the very real physical risks.

25. Are there IoT policy areas that could be appropriate for multistakeholder engagement, similar to the NTIA-run processes on privacy and cybersecurity?

A single large stakeholder organization, even on the order of the Networking and Information Technology Research and Development group, may prove inadequate. Specific communities of interest that have already been targeted by unique threats include those at risk for domestic violence (e.g., the StealthGenie app which was found to violate wiretapping statutes) and geofencing of patients of women's health services (e.g., Copley Advertising which identifies then targets women in Planned Parenthood).

A multistakeholder engagement that recognizes the scope of the problem and the need to include non-traditional stakeholders (e.g., patient advocates, caregivers of the elderly, domestic violence service providers) as well as technologists (from government, academic, and industry domains) can make progress on setting the requirements for privacy and security best practices in the IoT. A shared understanding of best practices in design for security and privacy in IoT is essential.

Summary

Computer security and privacy for an IoT ecosystem are both fundamentally important and fundamentally challenging. They are important because security and privacy lapses in IoT devices can cause real and significant harms to people and their environments. They are challenging because of the technical properties of IoT devices (the "technical element") and the complex issues that arise when designing technologies for a diverse collection of stakeholders (the "human element"). Given the diversity of challenges, we argue that any significant advance in the state of the art in security and privacy for an IoT ecosystem will require a large, interdisciplinary effort beyond engagement of traditional standards bodies to include a broadly defined set of stakeholders.

In considering effective security and privacy policies for IoT ecosystems, a challenge that we must overcome is that IoT devices have high diversity in terms of their technical capabilities and behaviors, as previously noted, and that IoT stakeholders have high diversity in terms of technical expertise. For example, some IoT devices may interconnect with other IoT devices, others may operate in isolation. Some IoT devices may have numerous sensors, others may have only a few or no sensors. Some IoT devices may directly affect their environment, others may only passively observe. Some IoT devices may receive data from or send data to the cloud, others may not. The interactions among the individuals, the devices, and the context of use create both a foundation for a threat model and the requirements for mitigating those threats.

The use of networked technologies that are aware, active, and adaptive create new security and privacy concerns. Artificial intelligence and continuous monitoring must allow for risk communication when people most need it, and individuals must be empowered to identify when discontinuities impinge safety. Individuals should be able to make explicit demands on the

technology, and receive timely responses. Issues of what should be done with the continuous dataflows that are an inherent part of the IoT (who should have access, when data should be deleted) are neither trivial nor self-evident. Multistakeholder processes and best practices are crucial.

References available upon request.

L Jean Camp, PhD; Ryan Henry, PhD; Steven Myers, PhD; Gianpaolo Russo, J.D. exp. 2017
School of Informatics and Computing
Indiana University
Bloomington, IN
{ljcamp, henry, samyers, russog} @indiana.edu