CENTER FOR LONG-TERM CYBERSECURITY

CLTC OCCASIONAL WHITE PAPER SERIES

Privacy and the Internet of Things

EMERGING FRAMEWORKS FOR POLICY AND DESIGN

GILAD ROSNER AND ERIN KENNEALLY

CLTC OCCASIONAL WHITE PAPER SERIES

Privacy and the Internet of Things

EMERGING FRAMEWORKS FOR POLICY AND DESIGN

GILAD ROSNER, PH.D.
Founder, IoT Privacy Forum

ERIN KENNEALLY, J.D.
Cyber Security Division, Science & Technology Directorate, U.S. Department of Homeland Security

(Author contributions to this work were done in her personal capacity. The views expressed are her own and do not necessarily represent the views of the Department of Homeland Security or the United States Government.)

CLTC
Center for Long-Term Cybersecurity

CENTER FOR LONG-TERM CYBERSECURITY

Contents

Executive Summary

The proliferation of network-connected devices, also known as the "Internet of Things" (IoT), offers unprecedented opportunities for consumers and businesses. Yet devices such as fitness trackers, personal home assistants (e.g., Amazon Echo, Google Home), and digital appliances are changing the nature of privacy as they operate silently in the background while transmitting data about a broad range of human activities and behaviors.

As "smart" becomes the new default setting for devices, consumers are further losing the ability to monitor and control the data collected about them, and they often have little awareness of what is done with their data downstream. The risks of sharing data through smart devices are not always clear, particularly as companies combine data from different sources to infer an individual's habits, movements, and even emotions.

This report provides an overview of some of the key privacy issues resulting from the expansion of the IoT, as well as emerging frameworks that could help policymakers and corporate leaders reduce potential harms through regulation and product design. Among the findings outlined in this paper:

- The IoT has the potential to diminish the sanctity of spaces that have long been considered private, and could have a "chilling effect" as people grow aware of the risk of surveillance. Yet the same methods of privacy preservation that work in the online world are not always practical or appropriate for the personal types of data collection that the IoT enables.

- Several frameworks have emerged for addressing the privacy issues that the IoT presents. Some focus on giving users more meaningful, granular control over the data that is collected, when data is collected, and how it is shared, while others focus on the accessibility and correct timing of privacy notices.

- Policymakers should take steps to regulate the privacy effects of IoT *before* mass sensor data collection becomes ubiquitous, rather than after. Omnibus privacy legislation can help regulate how data is handled in the grey areas between sectors and contexts. Europe's General Data Protection Regulation (GDPR), coming into force in 2018, will have an impact initially on IoT devices created and sold in the EU, and will affect those from the US as well over time.

- Having broad non-specialist conversations about the use, collection, and effects of IoT data is essential to help the populace understand technological changes in this space and how they affect privacy expectations.

- Makers of IoT products and services should employ a variety of standard measures to provide greater user management and control, as well as more effective notification about how personal data is captured, stored, analyzed, and shared.

The findings in this paper were developed through two workshops, seventeen semi-structured interviews, and an extensive literature review. A detailed analysis can be found in the full research report, *Clearly Opaque: Privacy Risks of the Internet of Things*,[1] which was funded by the William and Flora Hewlett Foundation. Sample quotations from these interviews and workshops are included throughout this paper.

Acknowledgments

The authors wish to thank all of the interviewees and workshop participants who contributed their valuable time and thinking to this research; a complete list can be found in the full-length report, *Clearly Opaque: Privacy Risks of the Internet of Things*. The authors also wish to thank Chuck Kapelke, Allison Davenport, and Betsy Cooper from the UC Berkeley Center for Long-Term Cybersecurity for their thoughtful edits, Eli Sugarman of the William and Flora Hewlett Foundation for funding this work and hosting a Bay Area workshop, Ian Wallace and the New America Foundation for hosting a workshop in Washington, D.C., and the numerous colleagues, practitioners, and scholars who reviewed drafts of the full report, whose names are listed therein.

Internet of Things 101

The Internet of Things emerged from a number of overlapping trends: widespread and inexpensive network access, cheap sensors and computing power, miniaturization, location positioning technology, inexpensive prototyping, and the ubiquity of smartphones as a platform for device interfaces. The connected devices of the IoT do *not* include multi-use computing platforms like laptops, tablets, or phones. Instead, they are products built for a narrow range of functions, and they share the ability to sense, analyze, and communicate.

Predictions vary widely about how many IoT devices are in the world and how many are coming. In 2012, IBM forecast there would be one trillion connected devices by 2015;[2] this did not come to pass. Cisco's widely used 2011 prediction anticipated 50 billion devices by 2020.[3] Gartner Research's oft-cited analysis claimed there were 8.4 billion devices in 2017, and they expect 20 billion in 2020.[4] Recently, a company called Statista predicted there will be 75 billion devices in 2025.[5] These numbers and their accompanying breathless predictions of market value should be taken with a grain of salt. As the IBM prediction illustrates, it's easy to get this wrong. It's also not always clear what the predictions refer to, as they vary in their inclusion of mobile phones and laptops, industrial devices, and IP-based and non-IP-based devices. As a result, the actual number of devices now and in the future is difficult to pinpoint.

The Internet of Things comprises an incredibly diverse range of products. To the right is a partial list, one that can keep growing because the IoT is a broad term that will ultimately encompass most digital products. If something can have a sensor and networking functions attached to it, it can be considered within the Internet of Things.

Consumer

- smart speakers
- connected cars
- intelligent door locks
- fitness and health wearables
- smart lighting
- networked thermostats
- smart TVs
- robot vacuums
- internet-connected toys
- networked bathroom appliances
- indoor security systems
- smart locks

Enterprise and Industrial

- worker productivity tracking devices
- smart office lighting
- temperature-sensitive supply chain
- augmented reality maintenance equipment
- autonomous trucking
- drones
- disease management systems
- employee wellness trackers
- automated retail checkout
- inventory optimization sensors
- face recognition cameras for security
- building management sensors

Regardless of shifting definitions and predictions, the IoT is still a useful concept for considering the economic, technological, and social impacts of a world of connected, sensing devices. Myriad reports, books, and articles have discussed how this evolution will benefit humanity. Many commercial organizations have highlighted the improvements and efficiencies gained by the introduction of smart devices, forecasting great benefits in the decades to come. Indeed, the IoT has the potential to improve road safety, free up time at home, improve health outcomes, make it easier to keep children safe, entertain us with richer experiences, make industrial processes cheaper and more efficient, help people conserve energy, and let us know ourselves better.

Yet these changes will result from the introduction of ever more sensors and computer processors into the human environment, including cameras, microphones, thermal sensors, motion detectors, facial and biometric analysis, identification technology, and environmental sensors. The introduction of such a broad and diverse sensor fabric into society has undoubted benefits, but it also introduces risks that must be explored and managed. This report focuses on the privacy risks that are emerging from the burgeoning Internet of Things, and examines how classic notions of private spaces are impacted by these sensing devices, how they affect people's ability to manage data about themselves, and what these devices means for society and our most cherished values.

Key Privacy Risks and Challenges

Despite the benefits that consumers will derive from IoT devices, there are also risks. One such risk is a change to how we see privacy. For the purposes of this report, privacy is defined as:

- the ability for people to *selectively share,* to determine how information about them is collected, used, and passed along;
- the ability to retreat from the gaze of and interactions with others;
- the right to be let alone, to create solitude and reserve from others;
- the ability to control the degree to which one is identifiable when undertaking online or offline activities; and
- the ability to control the *data impression* one gives off.

The Internet of Things heralds a qualitative shift in how privacy is managed, both by people and by the organizations that create, sell, and operate internet-connected devices. The IoT amplifies prior privacy challenges, such as the opacity of data flows and actors, and it creates new issues, such as enabling the stockpiling of emotional data. The sections below outline some of the key privacy risks and challenges related to the growth of the IoT, including the increase in online data collection, diminishment of private spaces, encroachment upon bodily and emotional privacy, challenges to meaningful consent, and regulatory issues.

A SHIFT FROM ONLINE TO OFFLINE DATA COLLECTION

Users of the internet share troves of information as they surf the web, including what web pages they visit, how long they spend on each page, and where they click on the screen. Through their behavior and voluntary sharing of data, they also frequently reveal personal information such as age, gender, income, and geographic location. This type of granular data collection has become so ubiquitous that it is expected, or met with resignation,[6] as a part of using the internet through a computer or mobile device.

As the Internet of Things expands, this type of granular data collection is moving into domains that have traditionally been considered "offline." The IoT enables an increase in monitoring of human activity that is fueled by *scale*—a greater number of sensing devices and sensor types—as well as a greater *proximity* of sensing devices to people's bodies and intimate spaces.

The commercial market offers devices that are intended to monitor people's activities and environments, as well as their physical bodies and emotions. In-home personal assistants, for example, bring always-on[7] microphones and cameras to spaces that were previously considered to be private, incorporating artificial intelligence and a melding of personal profile information gleaned from other sources. Health-tracking devices can transmit up-to-the-second details about a person's fitness, fertility, and heart health.[8] Nest, once a maker of smart thermostats, has expanded into indoor surveillance cameras.[9]

Even if a person does not invite these devices into their homes or onto their bodies, web-connected surveillance cameras, smart billboards, in-store retail tracking systems, and other public technologies are observing people's movements and habits on a massive scale.

Criminal exploitation represents an important concern with these devices, as each new bit of data stored represents a potential target for hackers. Yet the IoT also has potential to alter our lives in other ways, including by normalizing practices that in other contexts would be regarded as an invasion of privacy. The ultimate effects of this normalization are unclear: if children know that their teddy bear is watching them (or by extension, adults know that their smart TV is watching them), how will this affect their behavior? How do people meaningfully grant consent to be observed in a world of pervasive surveillance? How does the proliferation of internet-connected devices alter our traditional notions of privacy? And how should these devices be regulated to address these concerns?

DIMINISHMENT OF PRIVATE SPACES

Retreating to one's home, closing an office door, or hanging up a phone may have previously allowed a person to feel a measure of control over who might be listening or watching, but the presence of network-connected devices in private spaces can remove this sense of control and privacy. Experts warn that individuals' awareness of IoT devices' always-on technology can lead to chilling or conforming effects on behavior;[11] because these effects are difficult to quantify and study, such effects could go unnoticed or unaddressed.

> *"The IoT has the potential to really shift the home from a black box, what used to be a protective, safe space, to more of a glass house where everything that we do is now readily apparent to people who are willing to look for it."*
>
> *—Heather Patterson, Intel*[10]

From a regulatory perspective, connected devices pose problems for existing legal regimes such as the third-party doctrine,[12] which says that users give up their right to privacy when allowing third parties to collect and process their data, and the "reasonable expectation of privacy."[13] With the rise of ubiquitous data collection throughout the human environment, the notion of a private space may erode, and the ability to know who is observing us may cease to exist.

In addition to the "approved" uses of data, the IoT's massive collection of personal information creates a vast attack surface for malicious actors; indeed, the myriad sensors and actuators offer an opportunity to *weaponize* IoT to collect, use, and disclose data in ways that have a negative impact on privacy. There is a direct relationship between the IoT's technical underpinning—persistent and widespread collections and connectedness—and the likelihood that malicious actors will attempt to exploit sensitive personal information for economic gain. The potential for illicit use of data should be factored into all conversations about IoT privacy.

Legitimate vs Illegitimate Uses of Data

Privacy and security are related, often overlapping topics, though they have some fundamental differences. One concerns the distinction between *legitimate* and *illegitimate* uses of data. An illegitimate use of data is one that is unauthorized, i.e. when data is stolen, altered, or viewed by the wrong party. This is the domain of security, which protects data from being inappropriately accessed, modified, or shared. Legitimate uses of data are those that have been authorized.

However, in a discussion of privacy, there are plenty of legitimate data uses that may be problematic or harmful. For example, in countries where companies can collect individuals' data with only minimal notification, requiring users to search for ways to opt out, personal data can be used in ways that people did not expect or knowingly consent to. This is the domain of privacy, which is broadly concerned with how people control and manage data about themselves.

In essence, just because something is legal doesn't mean it is positive. While illegitimate uses of data must be combatted with security, legitimate but harmful uses of data must be interrogated through the lens of privacy preservation.

BODILY AND EMOTIONAL PRIVACY

The potential collapse of private spaces refers not only to physical spaces, but also to personal spaces, including our bodies. In the United States, laws exist that protect our bodies from certain types of collection; for example, a person cannot be forced to submit to a blood draw except in rare cases or with a warrant.[14] However, implantable chips, fertility trackers, and pills that can communicate are altering these boundaries. As wearable devices track bodily functions such as heart rate, temperature, and other data, people deserve to have clear understanding about *who* is collecting this data—and how they intend to use it.

> *"[F]irms can increasingly choose when to approach consumers, rather than wait until the consumer has decided to enter a market context. . . . In an age of constant 'screen time,' however, in which consumers carry or even wear devices that connect them to one or more companies, an offer is always an algorithm away. This trend of firms initiating the interaction with the consumer will only accelerate as our thermometers, appliances, glasses, watches, and other artifacts become networked into an 'Internet of Things.'"*
>
> *—Ryan Calo, "Digital Market Manipulation"*[15]

The IoT raises concerns about *emotional* privacy, as some connected devices have the ability to sense the emotional states of individuals through facial data, sentiment analysis, biomet-

rics, voice analysis, and other cues. Such technologies open the door to customized emotional manipulation for marketing or other purposes. Several industries have already indicated an interest in these emotional pictures of customers, including automobile firms, insurance providers, healthcare companies, recruitment agencies, advertising and marketing firms, and retail businesses.[16] The rise of emotion detection and "affective computing"[17] marks uncharted territory and calls for the establishment of new norms and regulations.

CHOICE AND MEANINGFUL CONSENT

In time, consumers may be unable to buy products that are not connected or that lack cameras and sensors. A reduced availability of "dumb" products versus "smart" ones can lead to an *erosion of choice,*[18] adding to the challenge of opting out of continual, ambient data collection.

> *"In the home environment you don't really have that much control over your privacy with IoT devices. Your biggest control element is deciding what device you place in your home and vetting them for good privacy practices. It's often difficult to find this information for consumer devices and take it into account in any kind of purchasing decisions."*
>
> *—Florian Schaub, University of Michigan School of Information*[19]

Even if consumers consent to the use of a device, whether they are *knowingly consenting*—i.e., understanding the full range of what they are sharing and how that data is used—is often unclear. Typically, consent for data collection by IoT devices operates on a "fire and forget it" basis: customers are presented with lengthy privacy policies up front and are given a binary choice to fully consent or not use the product. Following this initial agreement, consumers have little to no opportunity to withdraw consent.

The design of these products adds to the challenge: while computers and mobile devices have screen-based interfaces, IoT devices often lack screens, and so consumers cannot easily change privacy settings or access details about what data they are sharing. Research by experts reveals serious shortcomings in how products provide information about data collection and privacy to users.[20] Manufacturers are vague about what sensors are built into IoT devices and about which types of data constitute personal data. Few devices include a privacy policy in their physical packaging; instead, manufacturers provide links to websites, where the privacy policies are often difficult to find or insufficiently address privacy issues related to the device.

When users of a smart device are presented with a full privacy policy at the outset, these long, convoluted contracts often leave consumers with little understanding of what they are consenting to. Many companies that capture personal data are not even certain about what they will do with this data in the future, reducing users' ability to be fully informed about potential uses of collected data. The issues are thornier for devices designed for children (see sidebar). Children are not equipped to consent to data collection and use policies, so it is left up to parents to do so. The IoT, with weaker notifications about privacy and opaque chains of data collectors, makes it even harder for parents to protect their children's privacy. However, it is widely accepted that most people do not read privacy policies.[21] Parents risk making their children's play and behavior visible to many third parties, and neither they nor their children are likely to be aware of it.

The combination of a lack of screens and lax disclosure of privacy information makes it hard for purchasers of IoT products to understand what these devices see, hear, and know, as well as how their manufacturers and other parties will use the collected data.

A Case Study: Hello Barbie

To help conceptualize the privacy issues inherent to the IoT, consider the example of Hello Barbie, the first network-connected, interactive version of the classic Barbie doll. Released by Hasbro in 2015, this doll greets children with the phrase, "You're my best friend. I can tell you anything." Children can speak to Barbie by pressing a button on her belt: the audio files of the speaker's voice are encrypted and sent to an online speech analysis platform, which sends back an appropriate statement for the doll to respond.

Parents have access to their child's recordings, and a web-based interface makes it easy to share the recordings on social media. Do children understand that when they play with Barbie she's actually sharing their voice with other people? Do parents fully understand who all of the companies are that can access the recordings? Hello Barbie's maker gets things right by requiring a button to be pushed prior to recording and by encrypting all of the data, but questions about children's privacy still remain.[22]

From a security perspective, encryption is vital. Similar toys have experienced major privacy breaches in recent years. In February 2017, 2.2 million voice files from microphone-enabled teddy bears were compromised, and the related data was held for ransom.[23] In the same month, Germany banned a doll called "My Friend Cayla" that had such poor security that hackers on the other side of the world were able to take it over and speak through it.[24]

REGULATORY ISSUES SPECIFIC TO THE IOT

Regulating privacy in the IoT has many unique challenges. One issue is that internet-connected technologies often span multiple regulatory fields. For example, depending on a product's functionality and the data it collects, a single health tracking device might fall under the jurisdiction of the Department of Health and Human Services (HHS), the Federal Trade Commission (FTC), or the Food & Drug Administration (FDA). If an app is collecting health information, that collection and sharing may be governed by the HHS under the Health Insurance Portability and Accountability Act (HIPAA), but if the app makes recommendations about a person's health and wellness, it might come under the purview of the FDA.[25]

This muddling of jurisdictions means that a single device may have to adhere to several regulatory frameworks. The primary concern is not over regulation by multiple sectors, but an abdication of authority, as each agency passes the responsibility of enforcement to the others. For example, the National Highway Traffic Safety Administration released guidance on automated vehicles in September 2016 that included privacy guidance, including data minimization, but one year later removed all references in an update, saying instead, "the FTC is the chief Federal Agency charged with protecting consumers' privacy and personal information."[26]

In general, the question remains: does the IoT warrant its own regulations, or do existing policies suffice? In some ways, this question is mooted by U.S. states that are forging ahead with their own laws and policies, including those that regulate the privacy of vehicle event data recorder ("blackbox") information[27] or the privacy of imagery collected by drones flying over private and public spaces.[28] However, at the federal level, there is vigorous debate as to whether the IoT is deserving of new privacy regulations to address its new technical characteristics.[29]

Emerging Frameworks and Strategies

Our research revealed a variety of frameworks and approaches that could be useful for addressing questions about privacy and the IoT. Effective solutions will include a combination of governance regimes, adoption of strong standards within industry, and product design choices that prioritize user control and understanding.

OMNIBUS PRIVACY POLICY

> *"While US privacy protections are sectoral, data flows in the real world are not. As more objects get connected to the Internet, it will be more and more difficult to confine their data within a single regulatory silo."*
>
> *—Anna Slomovic, "Workplace Wellness, Privacy and the Internet of Things"*[30]

An omnibus privacy law has the potential to fill gaps left by ineffective or non-existent sectoral regulation and could improve the state of privacy not only for the IoT but arguably all internet technologies. A robust policy that encompasses all domains of personal data would give users more knowledge about what data is collected and more control over what is done with that data. A single regulatory framework would provide users and manufacturers with necessary clarity, and establish a better baseline for citizen's privacy expectations. Similarly, federal data security legislation would go a long way in ensuring that personal data is sufficiently protected by its custodians. Indeed, both the FTC and the Department of Commerce have been vocal about the need for such legislative protections. Despite this, omnibus federal privacy legislation has yet to reach an advanced stage in Congress, and the current administration's preference for deregulation reduces the already low chance of such legislation passing.

Europe's General Data Protection Regulation (GDPR) is a model for such an omnibus approach as it applies to all personal data, irrespective of type or the sector in which it was collected. The new law comes into effect in May 2018, and represents a substantial upgrade to the EU's existing omnibus data protection rules, the 1995 Data Protection Directive. The GDPR will affect American companies as the regulation applies to all entities that process Europeans' data, regardless of a company's geographic location. Compared to the existing US privacy regime, the GDPR requires far more internal assessment of data practices, and companies that fail to

comply face sanctions. This new framework will test the orthodoxy that increased regulatory burdens stifle innovation by corporations.

Unlike the United States' approach to privacy, requiring that a harm be shown to conclude that a privacy violation has occurred, the GDPR is oriented towards individuals' *rights*:

- the right to know how data about you is processed (collected, analyzed, and used)
- the right to object to such processing
- the right to see the data that is stored about you
- the right to a meaningful explanation about automatic data processing
- the right to withdraw consent to processing
- the right to have your data erased under certain conditions
- the right be able to easily move your data from one provider to a different one

The GDPR also requires data processors to maintain detailed records about the nature of their processing to be able to prove compliance to regulators. In most cases, companies that collect and process personal data will need to perform a data protection impact assessment (DPIA) to inventory the data they hold and determine how its processing affects the data subjects' rights. The GDPR also requires data processors to notify regulators about any data breaches without undue delay.

A regulation without effective enforcement mechanisms, however, would be toothless, and so the GDPR provides that companies could be fined up to 4% of their annual revenues for failing to comply. All of the GDPR's requirements, in combination with this sanctioning power, make it the most comprehensive data protection regulation in the world. It is also a way for Europe to 'export' its data protection and privacy norms to other parts of the world.

IMPROVED USER CONTROL AND MANAGEMENT

Manufacturers of IoT devices can help improve privacy standards by adopting practices or adding features that give users greater control over the data collected about them. All design elements should operate under the "least surprise principle": companies should be transparent and forthcoming, and not collect or use data in ways that violate people's expectations. Companies should commit to protecting users' privacy by only collecting data for which they have specific uses, versus hoarding it for some unknown, future use, and by deleting the data

when it is no longer needed. In addition, users should be given more power to update their privacy settings during the pre-collection or post-collection phases.

Companies should conduct *privacy impact assessments*, which help evaluate the impact and risks of collecting, using, and disseminating personally identifiable information. Privacy impact assessments are already mandatory for federal agencies in the U.S.—and for many companies in Europe, under the GDPR— and they have potential to help organizations identify risks, ensure compliance with laws, policies, or contracts, as well as put mitigation strategies in place.

To provide users with greater control, makers of IoT products should build in "Do Not Collect" switches or permissions, which would allow users to limit (or turn off) data collection. The most recognizable version of this is a "mute button" for devices with microphones. Companies can ensure their devices only begin data collection when a customer uses a "wake word" or manually activates collection. This is evident in devices like the Amazon Echo, which only starts to send spoken phrases to Amazon after someone wakes it up by saying, "Alexa." In general, products should indicate when they are monitoring people.

To improve the post-collection phase of data storage, companies can give users greater control by allowing them to withdraw consent to store data that has previously been collected. The GDPR requires that revoking consent must be as easy as granting it, an obligation that strongly supports user choice and control. Companies must also ensure that data is properly encrypted as it is transmitted—and after it has been received and stored—while giving users easy means to delete personal data.

Identity Management

Identity management (IDM) is the technical domain concerned with how people are identified within systems, how they authenticate to log in, who has authorization to see which information, and whether individuals can log in with pseudonyms or anonymous guest access. IDM is a valuable lens for considering the privacy posture of IoT devices, and offers useful concepts such as *unlinkability*, severing the links between users' activities on different devices, thus offering a narrower picture of their activities as a whole; and *unobservability*, making information about user activity invisible to intermediaries and transport networks. These two ideas should be incorporated into the design of IoT devices and platforms.

Different users of the same devices should be able to create separate profiles with different privacy settings, and have the option for pseudonymous use. Users should be able to easily switch between profiles and delete profiles that contain collected data. Devices with multiple users should separate profiles and their data collected from each user; one person should not be able to see the data of another person without explicit permission.

While discussions of privacy often focus on notions of hiding data from others, *selective sharing* is an essential privacy framing for the Internet of Things. The marriage of IoT devices and social networking allows people to share data from their fitness trackers, in-home devices, cars, toys, and other devices. But people don't want to share with everyone—they want to share this data *selectively* with appropriate parties (e.g., friends, fitness instructors, family member, doctors, etc.). *Privacy dashboards* and other similar design features can allow users to see, understand, and control the use and sharing of their personal data. Standards like the User-Managed Access (UMA) protocol enable developers to create a unified control point for users to authorize who can access their digital data, content, and services.[31]

Notification

In addition to building in design features that allow for greater user control, manufacturers can design devices to provide notifications to customers that are as transparent and as useful as possible.

The *timing* of a notice can have great effect on how well it communicates important information.[32] Privacy notices often appear during the setup of a device, and they tend to cover all current and future data collection over the lifetime of the product. However, other timing methods could be more effective, including:

- **Just-in-time notifications:** These appear just as data collection is about to occur so that a user can decide in real time if she wants to agree to sharing certain data.
- **Periodic notifications:** Regular reminders about ongoing data collection practices can allow users to reaffirm or cancel their consent at any time.

Example of periodic notification

- **Context-dependent notifications:** Notifications can be customized based on a user's context. For example, an alert about privacy risks could be sent when a user moves from inside to outside the home.
- **Layered notification:** This approach separates the granularity of notifications over time to give the user more information at the right time, and less when it's likely to be glossed over. For example, if a device's camera were not on by default, later, when the user decided to activate that feature, she would get a new privacy notice indicating that the device would now capture imagery and send it to the manufacturer for analysis. All four notification types can be used by a single device or service (see sidebar).

It is not enough, however, for manufacturers to simply update the timing of their notifications; they must also ensure that consumers *comprehend* these notices. Currently, privacy policies are not nearly as clear as they should be, as they are written by lawyers for lawyers. Product makers should conduct tests to determine whether users fully understand their data collection and use practices, and make improvements to their privacy policies based on user feedback. The chief method for this is to test comprehension with user groups prior to releasing a notice.

> *"Everything is stuffed into a privacy policy that of course no one reads, and they know that. I think it … illustrates how important trust truly is and the fact that companies do understand that, because they hide so much of what they're doing in either doublespeak or in these lengthy privacy policies or terms of service that they know that their users aren't likely to look at."*
>
> *—Michelle De Mooy, Center for Democracy & Technology*[33]

Researchers are also exploring how automation might enhance users' awareness of privacy aspects. For example, devices could be designed to automatically announce themselves so users are aware of their presence when entering spaces; device apps could also provide automatic

Augmented Notifications

Following is an example of just-in-time, periodic, layered, and context-dependent notifications:

- When a person uses a device feature she had not used before, she gets a notification about the types of data collected by that feature, explaining how it could be shared and what the privacy risks are. [just-in-time, layered]
- Once a month, the device reminds the user that it is collecting location information in the background, and displays a prompt for the user to affirm consent. [periodic]
- When the person is using a group feature (as opposed to using it solo), the device notifies her that data will be shared with the group. [just-in-time, context-dependent]

nudges to remind users about what data they are collecting.[34] Devices can also be designed to learn users' privacy preferences; for example: a notification could say, "You chose not to store GPS data when you are outside a one-mile radius of your home; would you also like to disable automatic check-ins at fitness facilities?"

Finally, regulators could support best practices in IoT governance by either requiring or nudging companies to design better notifications. For examples, regulators could:

- Provide guidance on best practices in notification in privacy policies;
- Require companies to collect feedback to assess consumers' comprehension of privacy policies;
- Expand the definition of personally-identifiable information to include data collected by IoT sensors;
- Require manufacturers to disclose what sensors are onboard devices and what they collect.

SUMMARY

Some of the frameworks and approaches above are more realistic or easier to implement than others. Legislation, for example, is a slow-moving process that is hard to influence without significant resources. However, for the makers of IoT devices and services, most of the suggestions for improving user control and management are reasonable and feasible:

- Design with the "least surprise" principle in mind
- Be maximally transparent about data collection and use
- Understand and stay within people's expectations
- Only collect data that has an immediate use, not a future, unspecified use
- Delete data as soon as it is no longer in use
- Perform a privacy impact assessment
- Ensure that products always indicate when monitoring is occurring

Improving the notifications provided to users should be easy for device makers, as these prompts can be either added at the inception of a new product or introduced after a device has been deployed through an update. Empowering consumers to manage their identities is admittedly more involved, as this must be considered in the early design phase of a system or platform.

Still, the technology marketplace is constantly pushing manufacturers to innovate at rapid speeds, and it only takes one or two product generations for significant changes to become widespread. The design suggestions provided above represent fruitful opportunities for companies that want to differentiate their products by providing users with more control over how their data is collected and used.

Conclusion

Most of the publicity around the Internet of Things has focused on cybersecurity risks, as media headlines have highlighted cases of hackers illegally accessing everyday products—such as cars, refrigerators, and children's toys—and using them for stealing data, spreading malware, or other nefarious purposes. Without doubt, industry leaders and regulators should invest significant time and resources in ensuring that all devices introduced to the IoT meet basic security protocols, such as encrypting data, requiring strong authentication, and automatically updating themselves with regular security updates.

At the same time, lawmakers and product designers should also ensure that, in addition to staving off hackers, IoT devices are designed to protect individuals' privacy *as part of their normal operation*, as the proximity and scale of IoT devices will collect people's activities, behaviors, and intimacies at an unprecedented scale. In this report, we detailed a variety of options available to implement robust frameworks to protect consumer privacy, whether by enabling greater user management and control, improving notification procedures, or advancing a robust policy framework. As the norms about when and where people expect to be observed shift and reasonable expectations of privacy evaporate, the laws related to these norms must be updated, and businesses should provide leadership in protecting consumer privacy.

Broad dialogue will be essential to help the public understand the nature of these technologies, particularly how they gather and share data. Rather than wait until privacy norms have already been eroded by the IoT, regulators and designers should work together now to build usable privacy into the products they create. Such measures will be essential to ensuring that our society continues to uphold the value of privacy as a fundamental right.

Endnotes

1 Gilad Rosner and Erin Kenneally, *Clearly Opaque: Privacy Risks of the Internet of Things* (London: Internet of Things Privacy Forum, 2018), https://www.iotprivacyforum.org/clearlyopaque.

2 Jon Iwata, "Making Markets: Smarter Planet," IBM Investor Briefing, May 9, 2012, https://www.ibm.com/investor/events/investor0512.html.

3 Dave Evans, "The Internet of Things: How the Next Evolution of the Internet Is Changing Everything," Cisco White Paper, 2011, https://www.cisco.com/c/dam/en_us/about/ac79/docs/innov/IoT_IBSG_0411FINAL.pdf.

4 Gartner, "Gartner Says 8.4 Billion Connected 'Things' Will Be in Use in 2017, Up 31 Percent From 2016" (press release, Gartner, February 7, 2017), https://www.gartner.com/newsroom/id/3598917.

5 Statista, "Internet of Things Devices Connected Devices Installed Base Worldwide From 2015 to 2025," *Statista*, accessed January 11, 2018, https://www.statista.com/statistics/471264/iot-number-of-connected-devices-worldwide/.

6 Joseph Turow, Michael Hennessey, and Nora Draper, "The Tradeoff Fallacy: How Marketers Are Misrepresenting American Consumers And Opening Them Up to Exploitation," Annenberg School of Communication White Paper, June 2015, https://www.asc.upenn.edu/sites/default/files/TradeoffFallacy_1.pdf.

7 For important nuance on the nature of "always-on," see Stacey Gray, "Always On: Privacy Implications of Microphone-Enabled Devices," Future of Privacy Forum White Paper, April 2016, https://fpf.org/wp-content/uploads/2016/04/FPF_Always_On_WP.pdf.

8 See, *e.g.*, the Ava fertility tracking bracelet, "Ava," accessed April 23, 2018, https://www.avawomen.com, and the QardioCore, a wearable ECG device, "QardioCore," accessed April 23, 2018, https://www.getqardio.com/qardiocore-wearable-ecg-ekg-monitor-iphone/.

9 Nest, "Nest Cam Indoor," accessed Jan 11, 2018, https://nest.com/cameras/nest-cam-indoor/overview/.

10 Heather Patterson (Intel), in discussion with the authors, August 2017.

11 Margot E. Kaminsky, "Robots in the Home: What Will We Have Agreed To?" *Idaho Law Review* 51, no. 3 (2015): 611.

12 Jennifer Lynch, "Will the Fourth Amendment Protect 21st-Century Data? The Court Confronts the Third-Party Doctrine." *SCOTUSblog*, August 2, 2017, http://www.scotusblog.com/2017/08/symposium-will-fourth-amendment-protect-21st-century-data-court-confronts-third-party-doctrine/.

13 Joel Reidenberg, "Privacy in Public," *University of Miami Law Review* 69, no. 1 (2014): 141.

14 David C. Sarnacki, "Analyzing the Reasonableness of Bodily Intrusions," *Marquette Law Review* 68, no. 1 (1984): 130.

15 Ryan Calo, "Digital Market Manipulation," *George Washington Law Review* 82, no. 4 (2014): 995.

16 Andy McStay, *Emotional AI: The Rise of Empathic Media* (London: Sage Publications, 2018).

17 R.W. Picard, "Affective Computing," *M.I.T. Media Laboratory Perceptual Computing Section Technical Report*, no. 321 (1995): 1.

18 Office of the Privacy Commissioner of Canada, "The Internet of Things: An Introduction to Privacy Issues with a Focus on the Retail and Home Environments," *Policy and Research Group of the Office of the Privacy Commissioner of Canada*, February 2016, https://www.priv.gc.ca/en/opc-actions-and-decisions/research/explore-privacy-research/2016/iot_201602/.

19 Florian Schaub (University of Michigan) in discussion with the authors, July 2017.

20 Scott Peppet, "Regulating the Internet of Things: First Steps Toward Managing Discrimination, Privacy, Security, and Consent," *Texas Law Review 93*, no. 1 (2014): 85; UK Information Commissioner's Office, "Privacy Regulators Study finds Internet of Things Shortfalls," *Information Commissioner's Office* (blog), September 22, 2016, https://ico.org.uk/about-the-ico/news-and-events/news-and-blogs/2016/09/privacy-regulators-study-finds-internet-of-things-shortfalls/.

21 Aleecia M. McDonald and Lorrie Faith Cranor, "The Cost of Reading Privacy Policies," *I/S: A Journal of Law and Policy for the Information Society* 4, no. 3 (2008): 543.

22 Meg Leta Jones and Kevin Meurer, "Can (and Should) Hello Barbie Keep a Secret?" 2016 IEEE International Symposium on Ethics in Engineering, Science and Technology, Vancouver, BC, May 13-14, 2016, https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2768507.

23 Troy Hunt, "Data From Connected CloudPets Teddy Bears Leaked and Ransomed, Exposing Kids' Voice Messages," *TroyHunt.com* (blog), February 28, 2017, https://www.troyhunt.com/data-from-connected-cloudpets-teddy-bears-leaked-and-ransomed-exposing-kids-voice-messages/.

24 Bill Chappell, "Banned in Germany: Kids' Doll is Labeled an Espionage Device," *NPR*, February 17, 2017, https://www.npr.org/sections/thetwo-way/2017/02/17/515775874/banned-in-germany-kids-doll-is-labeled-an-espionage-device.

25 US Department of Health and Human Services, "Examining Oversight of the Privacy & Security of Health Data Collected by Entities Not Regulated by HIPAA," US Department of Health and Human Services White Paper, 2016, https://www.healthit.gov/sites/default/files/non-covered_entities_report_june_17_2016.pdf.

26 "Automated Driving Systems: FAQ: Voluntary Guidance," *US National Highway Traffic Safety Administration*, accessed January 11, 2018, https://www.nhtsa.gov/manufacturers/automated-driving-systems.

27 "Privacy of Data from Event Data Recorders: State Statutes," *National Conference of State Legislature*s, January 29, 2017, http://www.ncsl.org/research/telecommunications-and-information-technology/privacy-of-data-from-event-data-recorders.aspx.

28 National Conference of State Legislatures, "2017 Unmanned Aircraft System State Legislation Update," January 17, 2018, http://www.ncsl.org/research/transportation/2017-unmanned-aircraft-systems-uas-state-legislation-update.aspx#privacy.

29 See, *e.g.*, US Department of Commerce, "Fostering the Advancement of the Internet of Things," US Department of Commerce Green Paper, January 12, 2017, https://www.ntia.doc.gov/other-publication/2017/green-paper-fostering-advancement-internet-things.

30 Anna Slomovic, "Workplace Wellness, Privacy and the Internet of Things." *AnnaSlomovic.com* (blog), January 31, 2015, https://www.annaslomovic.com/single-post/2015/01/31/Workplace-Wellness-Privacy-and-the-Internet-of-Things.

31 "Workgroup – User Managed Access," Kantara Initiative, March 3, 2018, http://tinyurl.com/umawg.

32 Florian Schaub et al., "A Design Space for Effective Privacy Notices," Symposium on Usable Privacy and Security (SOUPS) 2015, Ottawa, Canada, July 22–24, 2015.

33 Michelle De Mooy (Center for Democracy & Technology) in discussion with the authors, August 2017.

34 See, *e.g.*, "The Personalized Privacy Assistant Project," accessed April 23, 2018 https://www.privacyassistant.org/.

About the Authors

Dr. Gilad Rosner is a privacy and information policy researcher and the founder of the non-profit Internet of Things Privacy Forum, whose mission is to produce guidance, analysis, and best practices to help industry and government reduce privacy risk and innovate responsibly in the domain of connected devices. Gilad's broader work focuses on identity management, US & EU privacy and data protection regimes, consumer protection, and public policy. His research has been used by the UK House of Commons Science and Technology Committee report on the Responsible Use of Data and he has contributed directly to US state legislation on law enforcement access to location data, access to digital assets upon death, and the collection of student biometrics. Gilad has consulted on trust issues for the UK government's identity assurance program, Verify.gov, and is the author of *Privacy and the Internet of Things* (O'Reilly Media, 2017).

Gilad has a 20-year career in IT, having worked with identity technologies, digital media, automation, and telecommunications. Prior to becoming a researcher, he helped design, prototype, and manufacture the world's only robotic video migration system, known as SAMMA, which won an Emmy Award for technical and engineering excellence in 2011. Gilad is a member of the UK Cabinet Office Privacy and Consumer Advisory Group, which provides independent analysis and guidance on Government digital initiatives, and a member of the IEEE Global Initiative for Ethical Considerations in Artificial Intelligence and Autonomous Systems. He is a Visiting Scholar at the UC Berkeley School of Information and a Visiting Researcher at the Horizon Digital Economy Research Institute. Gilad is a graduate of the University of Chicago, NYU, and the University of Nottingham.

Erin Kenneally is a Program Manager in the Cyber Security Division for the Homeland Security Advanced Research Projects Agency (HSARPA) at DHS S&T. Her portfolio includes trusted data sharing, privacy and Information Communication Technology (ICT) ethics, including managing the IMPACT (Information Marketplace for Policy and Analysis of Cyber-risk and Trust), cyber risk economics (CYRIE) and data privacy projects. Prior to joining CSD, Kenneally was Founder and CEO of Elchemy, Inc., and served as Technology-Law Specialist at the International Computer Science Institute (ICSI) and the Center for Internet Data Analysis (CAIDA) and Center for Evidence-based Security Research (CESR) at the University of California, San Diego.

Erin is a licensed attorney specializing in information technology law, including privacy technology, data protection, trusted information sharing, technology policy, cybercrime, ICT ethics, and emergent IT legal risks. She holds Juris Doctorate and Masters of Forensic Sciences degrees, and is a graduate of Syracuse University and The George Washington University.

CLTC

Center for Long-Term Cybersecurity

UC Berkeley

Center for Long-Term Cybersecurity
cltc.berkeley.edu
@CLTCBerkeley