An Updated Privacy Paradigm for the "Internet of Things"

By Christopher Wolf and Jules Polonetsky
Co-Chairs, Future of Privacy Forum

November 19, 2013

The Future of Privacy Forum is a think tank whose mission is to advance privacy for people in practical ways that allow for innovation and responsible uses of data. The FPF Advisory Board includes representatives of business, privacy scholars and consumer advocates. www.futureofprivacy.org

Introduction

The "Internet of Things" refers to the information networks comprised of sensors and other technologies embedded in physical objects and linked via wired and wireless networks. Cisco estimates that there are nearly 11 billion connected objects in the world.[1] By 2020, there may be more than 200 billion connected devices.[2] As the Internet of Things matures, more and more everyday objects will "wake up," become aware of their environments, communicate the information that they collect, and receive information from outside sources. This will likely generate substantial economic and social benefits, including improved health care, increased public and personal safety,

[1] *Connections Counter: The Internet of Everything in Motion*, Cisco Newsroom, http://newsroom.cisco.com/feature-content?articleId=1208342 (last visited Oct. 29, 2013).

[2] *See* Press Release, International Data Corporation, *The Internet of Things Is Poised to Change Everything, Says IDC* (Oct. 3, 2013), *available at* http://www.idc.com/getdoc.jsp?containerId=prUS24366813.

efficient use of resources, business innovations, and more. This paper examines the need for an updated, forward-looking privacy paradigm for the Internet of Things.

The Current Privacy Paradigm Is Not Practical for the Internet of Things

Along with these potential benefits, the Internet of Things gives rise to debate over privacy and security concerns. In some cases, existing privacy concerns are heightened by the increased data interaction with newly interconnected objects. Therefore, the question is: *How do we account for privacy in the Internet of Things?*

Traditionally, privacy concerns have been addressed by application of the Fair Information Practice Principles ("FIPPs"), which address the treatment of personal information. In 1973, the United States Department of Health, Education, and Welfare offered the first comprehensive articulation of the FIPPs.[3] The FIPPs have since been embodied in U.S. and European Union privacy laws and serve as the basis for a range of privacy frameworks established by legislatures, government agencies, and international bodies.[4] The Organization for Economic Cooperation and Development developed one of the more influential variations of the FIPPs in its Guidelines on the Protection of Privacy and Transborder Flows of Personal Data ("Guidelines").[5] The Guidelines, which were adopted in 1980 and revised this year, are intended to "address concerns arising from the increased use of personal data and the risk to global

[3] FTC, Privacy Online: A Report to Congress 48 n.27 (1998), *available at* http://www.ftc.gov/reports/privacy3/priv-23a.pdf.

[4] *See e.g.*, The White House, Consumer Data Privacy in a Networked World: A Framework for Protecting Privacy and Promoting Innovation in the Global Digital Economy (2012); FTC, Protecting Consumer Privacy in an Era of Rapid Change 22 (2012). *See generally* John W. Kropf, *Independence Day: How to Move the Global Privacy Dialogue Forward*, Bloomberg BNA Privacy & Security Law Report (Jan 12, 2009)

[5] *See* Kropf, *supra* note 4.

economies resulting from restrictions to the flow of information across boundaries."[6] Since then, the FIPPs have been presented in different ways with different emphases.[7]

In their various formulations, the FIPPs establish core principles guiding the collection, use, and disclosure of data.[8] Some of the more important FIPPs are 1) Notice- individuals should be provided with timely notice of how their data will be collected, used, and disclosed; 2) Choice- individuals should be given choices about whether and how their data will be used; 3) Data Minimization- organizations should seek to limit the amount of personal data they collect and that might be retained; 4) Purpose Specification- the purposes for which personal data are collected should be specified prior to or at the time of collection; and 5) Use Limitation- personal data should only be used for those purposes specified prior to or at the time of collection.[9]

Privacy Challenges Presented by the Internet of Things Cannot be Solved by Simple Application of the FIPPs.

The FIPPs generally have been thought of as establishing high-level guidelines for the implementation of specific codes of practice.[10] They do not establish a specific set of rules prescribing how organizations must go about promoting privacy in all contexts. The FIPPs of notice and choice are often implemented in ways that are not well-suited for the Internet of Things, such as through the posting of privacy policies and the use of

---

6 OECD, OECD Guidelines Governing the Protection of Privacy and Transborder Flows of Personal Data 19 (2013), *available at* http://www.oecd.org/sti/ieconomy/2013-oecd-privacy-guidelines.pdf.

7 *See supra* note 4; Edith Ramirez, The Privacy Challenges of Big Data: A View from the Lifeguard's Chair, Keynote Address by FTC Chairwoman Edith Ramirez, Technology Policy Institute Aspen Forum (Aug. 19, 2013), *available at* http://www.ftc.gov/speeches/ramirez/130819bigdataaspen.pdf.

8 *See, e.g.*, *id.* at 13-14; The White House, *supra* note 4, at 10.

9 *See, e.g.*, OECD, *supra* note 6, at 14; The White House, *supra* note 4, at 11-19, 21.

10 *E.g.,* The White House, *supra* note 4, at 16 n.21.

click-through consent mechanisms. Many connected devices, such as traffic sensors embedded in roadways, will not be equipped with interactive screens or other user

interfaces. When the Internet of Things matures, it is likely that most connected devices will be invisible to us (*i.e.*, we will not interact directly with them frequently, if at all). Moreover, the individual owning or registering a device may lend that device to others. In those situations, the person operating the device may not have had the opportunity to provide consent to data collection. Although technological solutions may be developed to facilitate notice and choice options, it would be impractical to premise data collection and use in the Internet of Things on traditional notice and choice implementations.

The Internet of Things relies on frequent, often continuous, data inputs and transmissions from a broad array of connected devices. If the only way to authorize the collection of personal data were based on traditional notice and choice, individuals would be prompted to consent to data collection and use each time they bumped into new connected devices. That could occur hundreds or thousands of times a day. Not only would that substantially slow the data transmissions underlying the Internet of Things, it would be incredibly burdensome for individuals and could hinder the development of innovative new technologies.

It is unrealistic to expect that individuals will be willing or able to effectively register their informed preferences in a world where they are regularly prompted to read and accept notices of complex data collection, use, and sharing practices. Individuals may end up blindly accepting data practices rather than having to endure reading yet one more

privacy disclosure.[11] Instead of protecting privacy, strict adherence to traditional notice and choice principles may drive individuals to give up.

Purpose specification, data minimization, and use limitation also present problems for the Internet of Things. As mentioned before, those principles require organizations to specify the purposes for which they will use the data they collect, collect only that data needed to achieve those ends, and use the data only for specified purposes. That risks unduly limiting the development of new services and the discoveries that may follow from valuable research.

Consider the innovations pioneered by the United Nations Global Pulse that are enabled by the analysis of mobile phone data. Global Pulse has helped us understand mobility, social interaction and economic activity.[12] By analyzing mobile interactions, UN researchers were able to examine the post-earthquake population migration caused by the Haiti earthquake.

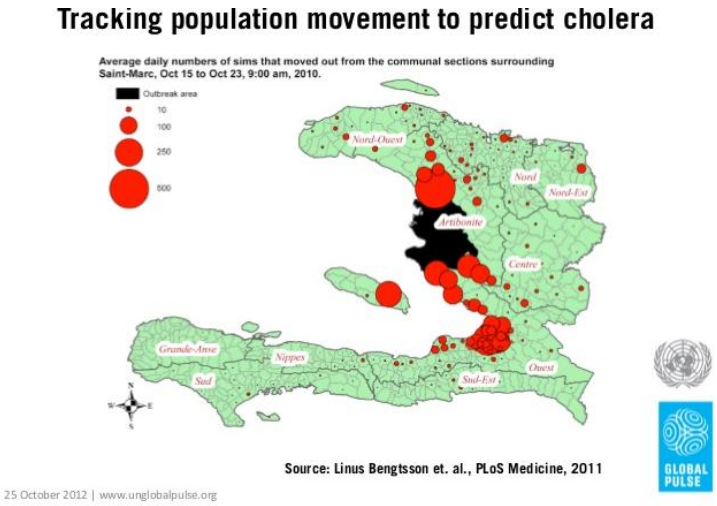

[11] *See generally* Fred H. Cate & Viktor Mayer-Schönberger, Notice and Consent in a World of Big Data, Int'l Data Privacy Law, Vol. 3, No. 2 (2013).

[12] Robert Kirkpatrick, Beyond Targeted Ads: Big Data for a Better World (2012), *available at* http://www.slideshare.net/unglobalpulse/strata-14934034.

Global Pulse has been able to map the areas in Kenya where Malaria was likely to spread and assess how well Mexico was combating the H1N1 virus. They were also able to better understand socio-economic activity in a number of countries, as well as to help plan road infrastructure and analyze traffic patterns.

Across the US, utilities have installed smart meters, seeking to help residents manage power use more effectively and benefit the environment. Utilities will be able to learn how to adapt and manage their systems—thereby securing the stability and efficiency of the smart grid—only by understanding how residents change their usage patterns. As electric vehicles increasingly are charged at home, understanding how and when drivers come home and plug-in their vehicles will be needed to ensure that the grid can adapt to changing patterns, lest we risk overburdening the system at the end of each evening commute. In the course of managing the smart grid, we are likely to uncover a host of surprising uses for which we can use data about power usage. We may discover that that data can be used to promote health or identify the need for new transportation, entertainment, or food storage technologies.

You cannot specify what you cannot imagine. If data can be processed only in accord with specified purposes, we risk losing out on the unimagined possibilities that the Internet of Things may provide. Our challenge is to allow practices that will support progress, and provide appropriate controls over those practices that should be forestalled or constrained by appropriate consent.[13]

[13] For example, determining the balance between the benefits of new uses and the attendant risks may in some instances require more sophisticated privacy impact assessments that can analyze the impact of risks or harms and assess the potential benefits for individuals and society. *See* Jules Polonetsky & Omer

The inadequacy of traditional privacy practices in the Internet of Things era is not entirely surprising. We tend to view privacy as "the claim of individuals, groups, or institutions to determine for themselves when, how, and to what extent information about them is communicated to others."[14] But the revolutionary impact of the Internet of Things derives largely from its reliance on myriad and continuous communications. To ask individuals to protect their privacy by managing those communications would be akin to telling Sisyphus that he can rest as soon as he gets that rock to settle atop the hill.

A Use-Focused Privacy Paradigm Is Well-Suited for the Internet of Things

Rather than focusing on how information is collected and communicated, we should rely on how personally identifiable information is used. The following proposals reflect how this can be done.

Use anonymized data when practical. When organizations use adequately anonymized data sets, their use of that data should not be restricted under privacy laws or regulations. As noted by the Federal Trade Commission in its 2012 privacy report, further privacy assurances can be obtained when organizations publicly commit to not re-identify data and when organizations contractually require the third parties to which they send anonymized data to not attempt re-identification.[15] Anonymizing personal information decreases the risks that personally identifiable information will be used for

---

Tene, Privacy and Big Data: Making Ends Meet, 66 Stan. L. Rev. Online 25, 26-27 (2013), http://www.stanfordlawreview.org/online/privacy-and-big-data/privacy-and-big-data.

[14] Alan F. Westin, Privacy and Freedom 7 (1967).

[15] FTC, Protecting Consumer Privacy in an Era of Rapid Change 22 (2012).

unauthorized, malicious, or otherwise harmful purposes.[16] Properly anonymized data are highly unlikely to have any impact on individuals and do not implicate privacy concerns.

Although there have been some reports of researchers who were able to re-identify information from supposedly anonymized data sets, it would be a mistake, however, to conclude that it is always easy to re-identify data or that anonymization is not a useful, privacy-protective practice. In 2009, a group of experts attempted to re-identify approximately 15,000 patient records that had been de-identified under the standards of the Health Insurance Portability and Accountability Act ("HIPAA"). They used commercial data sources to re-identify the data and were able to identify only .013% of the individuals.[17] When data sets are anonymized properly, re-identification is no easy task.[18] When anonymized data sets are kept securely in house with a strong commitment and internal checks to prevent re-identifying the data, then anonymization serves as a strong protection to address privacy concerns.

Whether a specific anonymization practice is appropriate will depend on the circumstances.[19] When anonymizing data, organizations should assess the risks that the data could be re-identified given the nature of the data, the context in which the data will be used, and the resources available to those with access to the data.

---

[16] *See* Ann Cavoukian & Khaled El Emam, Dispelling the Myths Surrounding De-identification: Anonymization Remains a Strong Tool for Protecting Privacy 4 (2011).

[17] Deborah Lafkey, *The Safe Harbor Method of De-Identification: An Empirical Test*, ONC Presentation, October 8, 2009, *available at* http://www.ehcca.com/presentations/HIPAAWest4/lafky_2.pdf.

[18] Cavoukian & El Emam, *supra* note 16, at 7.

[19] For example, gender is not an identifying characteristic if every member of a large data set is female. However, if there is only one woman in data set, using gender in the data set will facilitate identification of her.

Organizations that are inexperienced with anonymization should consider implementing third-party testing to determine the likelihood of re-identification.

With robust anonymization practices in place, organizations will be able to use information as needed to realize the mature development of the Internet of Things and spur tomorrow's headline technologies while promoting individual privacy.

Respect the context in which personally identifiable information is collected. This principle is often interpreted to mean that personally identifiable information should be used only in the ways that individuals would expect given the context of the collection. Consumers expect that companies will share personally identifiable information with other companies to fulfill orders and that companies will use personal information to engage in first-party marketing. When personally identifiable information is used in those ways or in others that individuals would reasonably expect, there is no privacy violation.

However, respect for context should not focus solely on what individuals would expect; there may be unexpected new uses that turn out to be valuable societal advances or important new ways to use a product or service. Consider a company that collects personal fitness information from wearable sensors that track sleep, steps taken, pulse or weight. Analysis of such data, collected originally only to report basic details back to users, may yield unanticipated health insights that could be provided individually to users or used in the aggregate to advance medical knowledge. Rigidly and narrowly specifying context could trap knowledge that is available and critical to progress.

Be transparent about data use. To complement the respect for context principle, organizations should be transparent about the purposes for which they will use personally identifiable information. Even if organizations cannot predict how they will use personally identifiable information in the Internet of Things, they can inform individuals that they will use such information to improve products, conduct research, or increase security measures. Organizations making decisions that affect individuals could, subject to protecting their intellectual property, disclose the high-level criteria used when making those decisions. Insurance companies, for instance, could disclose that they determine premiums solely by reviewing driving habits, location, driving history, and other permissible data categories. The insurance companies could clarify that factors such as ethnicity, sexual orientation, and political preferences are not factored into premium determinations.

The required levels of transparency and limitations to which data may be used in a given context should, however, be tailored to the level of identifiability of data, with adequately anonymized data being subject to fewer limits or restrictions.

Automated accountability mechanisms can be designed to determine how personally identifiable information is used and whether the uses conform to established policies. As data flows become more and more complex, it will become more and more difficult for individuals to monitor and enforce privacy compliance. To support privacy compliance, organizations should develop and implement automated systems that can monitor and assess the myriad uses and transmissions of personally identifiable information. Professor Hal Abelson at MIT has proposed that information be tagged with its provenance and logs of transfers and uses. Automated accountability mechanisms

could monitor data usage and determine whether the uses comply with machine readable policies.[20] When improper uses are identified (*e.g.*, credit is denied after viewing someone's political affiliation), accountability mechanisms could notify appropriate parties and trigger appropriate actions.

Develop Codes of Conduct. As the Internet of Things becomes more ubiquitous, parents will want to control what can be done with information collected from devices associated with their children. Others may want to indicate their preferences about how third-party connected devices will communicate with them. Self-regulatory codes of conduct will be the most effective means to honor these preferences and others in the rapidly evolving landscape of the Internet of Things. Codes of conduct could establish frameworks that enable individuals to associate usage preferences with their connected devices. These preferences would indicate to other devices how information collected from individuals' devices may be used. Preferences could serve as inputs for the accountability mechanisms discussed above, and robust codes of conduct (perhaps supported by audits of accountability mechanisms) could serve to establish accountability.

It's not too early to start, as FPF has pioneered codes of conduct for smart home devices and for smart stores and is coordinating a working group of connected car leaders.

Provide individuals with reasonable access to personally identifiable information. Businesses and other organizations could allow individuals reasonable access to and

[20] Hal Abelson, Information Accountability as the Foundation of 21st Century Privacy Protection (2013), *available at* http://kit.mit.edu/sites/default/files/documents/Abelson_MIT_KIT_2013_Conference.pdf.

use of their personally identifiable information. This will likely enhance consumer engagement with and support of the Internet of Things. One way to provide reasonable access would be to offer tools that allow users to add, tailor, or featurize data, perhaps by allowing access via third-party application programming interfaces. The more effectively that data is anonymized, the less the need and the ability to provide detailed access.

Conclusion

Time tested privacy principles will continue to have relevance for the Internet of Things, but policymakers will need to be flexible and creative in applying these principles to new technologies. As they evaluate their role, and that of industry, in protecting personal privacy and ensuring data security in the world of the Internet of Things, a rigid application of the current privacy paradigm is not practical or appropriate. Thus, we respectfully urge consideration of the updated privacy paradigm we have proposed for the Internet of Things.