A Digital Identity Stack to Improve Privacy in the IoT

Stephen Wilson, Nour Moustafa, and Elena Sitnikova
University of New South Wales
Canberra, Australia
s.wilson@student.adfa.edu.au nour.moustafa@unsw.edu.au, e.sitnikova@adfa.edu.au

***Abstract*— The Internet of Things increasingly involves collection, processing and transmission of a wide variety of data to services and other devices. Business and engineering considerations are both increasing the volumes and detail of IoT data flows. Reasonably obvious privacy risks result from IoT-connected devices when they emit identifiable information, for this can reveal the activities of device users. More subtle risks arise when bulk device data is available for analysis, and linkage to auxiliary data sets, because identification or re-identification of users can follow. At the same time, security engineers are now designing for the "Identity of Things", exploiting embedded cryptography and SIM-like modules to help with the authentication and authorization of devices acting as independent agents in the IoT. To help protect privacy while allowing precise authentication, this paper sets out a new model for digital identity management, comprising a stack of identities, attributes, and attribute metadata. As with the familiar OSI network stack, the digital identity stack helps to decouple different layers of authentication technology, so that IoT data is shared on an explicit need-to-know basis, and extraneous disclosures are minimized.**

***Index Terms*—Internet of Things, authentication, security, privacy, digital signatures, public key infrastructure.**

I. Introduction

The Internet of Things (IoT) features a wide range of instruments, consumer devices and regular domestic appliances, communicating automatically with each other and with network services, largely in order to streamline operations over near-ubiquitous communications and cloud computing infrastructure [1]. These sorts of devices are often labeled "smart" or "intelligent", as they use autonomous sensors and applications to direct real-time actions and interactions with other systems in the environment [2].

The value-added processing of raw data opened up by IoT and Big Data has led to comparisons between data and crude oil [3]. While academics and policy makers are increasingly aware of how data mining creates new privacy risks through identification and re-identification, regular users may be oblivious to how manufacturers refine raw data collected via their devices, and capitalise on the resulting information assets. Auto manufacturers for instance reserve the right to retain a variety of data including drivers' radio listening histories [4]. There is an expectation amongst some IoT entrepreneurs that "Automakers will make more money selling vehicle data than the cars themselves" [5].

Identity Management (IdM) in IoT has been framed by some in terms of the "identity of things". Most Internet-connected devices will become uniquely identifiable. This can make machines proxies for their users, going everywhere the users go, thus revealing where people go, when and how they go, and often who they go with. Services like the road traffic display in Google Maps monitor the location and movement of mobile phones determined to be travelling in vehicles. The raw data also carry strong identification of drivers and passengers in the form of mobile phone numbers and identifiers.

This paper discusses how digital identity can be systematized into a formal stack-like model (analogous to the OSI networking stack) as an aid to design and specifically to help IoT privacy. The model facilitates IoT "privacy by design" patterns, so that generally less personal information needs to be collected and transmitted, and leakage of identifying data avoided.

II. The privacy impact of IoT

In over 100 countries, legal privacy protection is given to a class of data known as *personal information* (PI), *personal data*, or *personally identifiable information* [6]. In general, data protection aka privacy laws generally apply only to PI. Most jurisdictions have broadly equivalent definitions. In the United Kingdom, PI is "data which relate to a living individual who can be identified" [7]; in Australia, it is "information or an opinion ... about an identified individual, or an individual who is reasonably identifiable" [8].

Engineers need to be aware that PI is generally defined in terms of potential identifiability; data may be cautiously classified as personal before they are *actually* identified.

Privacy legislation typically enshrines principles similar to those of the OECD. For our purposes, we highlight the following:

- Collection Limitation: "There should be limits to the collection of personal data and any such data should be obtained by lawful and fair means and, where appropriate, with the knowledge or consent of the data subject".
- Purpose Specification: "The purposes for which personal data are collected should be specified not later than at the time of data collection and the subsequent use limited to the fulfilment of those purposes or such others as are not incompatible with those purposes and as are specified on each occasion of change of purpose".

- Use Limitation: "Personal data should not be disclosed, made available or otherwise used for purposes other than those specified in accordance with the [Purpose Specification Principle]".
- Openness: "There should be a general policy of openness about developments, practices and policies with respect to personal data" [9].

A. Privacy risks and countermeasures in IoT

The overt privacy risks posed by the IoT have been catalogued by the Internet of Things Privacy Forum as follows [10]:

1. Enhanced Monitoring of individuals via connected devices
2. Nonconsensual Capture of PI by devices in ways often not understood
3. Collecting Medical Information from new health and fitness products
4. Breakdown of Informational Contexts
5. Diversification of Stakeholders, where appliance manufacturers become involved in information related business, and
6. Backdoor Government Surveillance.

Engineers working in the IoT need to be aware that, because data from multiple sensors may be combined downstream from the original collection to create new insights, the upstream sensor data on a precautionary basis should be regarded as personal information. De-identifying sensor data by stripping it of overt identifiers is a necessary but not sufficient privacy protection. Removing identifiers may render isolated data anonymous, yet Big Data processes allow disparate records to be linked, analysed, and often re-identified, in the absence of explicit identifiers.

So data analytics exacerbates the above risks in several ways. It is not merely the nonconsensual capture of PI that matters but the nonconsensual *synthesis* of fresh PI from raw data. More than data "mining", Big Data practices and business models are creating complex information supply chains where PI may be synthesised well downstream from IoT sensors. Similarly, there is great potential for medical information to be created as well as collected when raw data is processed (as when for example department stores use subtle changes in purchasing patterns to predict when customers have become pregnant). A key element of mitigating these sorts of risks is to manage the risk of identification or reidentification of IoT users, and this requires close attention to the relationship between digital identity and attributes, as discussed below.

B. Identification in the IoT

The privacy risks resulting from multiple IoT sensors and analytics have been summarised by Peppet: "each type of consumer sensor … can be used for many purposes beyond that particular sensor's original use or context, particularly in combination with data from other Internet of Things devices. Soon we may discover that we can infer whether you are a good credit risk or likely to be a good employee from driving data, fitness data, home energy use, or your smartphone's sensor data" [11].

Researchers are constantly finding ways to deduce identity from raw data. For example, diagnostic speed and direction data routinely logged by automotive black boxes can be used to compute drivers' whereabouts without absolute GPS coordinates [12]. And the identity of cell phone users can be disambiguated through just four data points comprising location of the phone (deduced from cell tower metadata) and the time of call [13]. So it is widely recognized that individuals in IoT can be indirectly identified by many means and exposed to privacy risks. Privacy protection in IoT requires further attention.

III. Developments in Identity Attributes

An archetypal definition of digital identity was provided by the so-called "Laws of Identity" developed by Microsoft, namely "a set of claims" about a digital entity, such as a user or a device [14]. Several IdM industry initiatives in recent years explicitly concentrate on attributes.

- The FIDO Alliance (https://fidoalliance.org) is developing technical protocols for authenticating users and mobile devices with reduced usage of passwords. A special feature of the FIDO protocols is the Metadata Service which publishes precise low-level statements about the state of a device [15]. Most messages sent from FIDO-enabled devices are digitally signed by cryptographic keys carried securely within the device. Most FIDO applications presently use mobile phones, but the authentication protocols are well suited to the IoT where embedded cryptographic processing should be widespread.
- The Vectors of Trust Internet Draft of the Internet Engineering Task Force (IETF) seeks to better characterize the reliability of a user's authenticators by gauging various independent contributing factors of identity and authentication processes, hence the reference to "vectors" [16].
- The Verifiable Claims working group of the World Wide Web Consortium (W3C) is working on protocols for expressing and exchanging dependable claims about users (such as qualifications, account details, roles and so on) [17].

Building on this increasing attention to low level attributes, the new stack model presented in this paper can help to better systematize digital identity analysis and design.

IV. A Proposed Digital Identity Stack

It has often been posited that the Internet needs an "identity layer" [18][19]. But the idea has lacked precision, and has not been explained in relation to other well-established layers in the network technology stack [20]. Given the practical importance of attributes and the maturity of authentication protocols, IdM elements can be segmented into a distinct new stack involving more than just a single identity layer. Figure 1 depicts the stack, which comprises the following layers:

- **Relationships**: At the top of the stack are relationships between users and service providers. Individuals commonly deal with multiple service providers including, in the context of the IoT, appliance manufacturers, service personnel, home builders, after-market software providers, insurers, finance providers, hire-purchase companies, utilities companies, certifiers, and product safety regulators.
- **Identities**: Online relationships are typically indexed by identities, often names, reference numbers or other identifiers. Each identity at this level of the stack is a handle by which users are known by vendors.
- **Attributes** are the various pieces of information service providers needs to know about the people they deal with. Attributes can be used in identification (in Know Your Customer processes for instance). Discrete attributes also help match services to users, without overt identification; for example, an IoT device may behave differently depending on what it knows about a user's driving license, home location, medical conditions, employment, insurance and so on.
- **Authentication Data**: at this layer, codified attributes are exchanged between sub-systems (typically client and server) using technical protocols such as FIDO Universal Authentication Framework (UAF), FIDO Universal Second Factor (U2F), OAuth (Open Authorization), Open ID Connect, or Security Assertion Markup Language (SAML).
- **Authentication Metadata**: additional 'data about data' are used to confer the quality or provenance of the attributes and authentication data; for example, a service may wish to know the source or issuer of an attribute, its age or expiry date, intended usage of the attributed, restricted uses, detailed policy information, and revocation status. Typically, authentication data and metadata are digitally signed, to bind them together and preserve their integrity.
- **Deeper Network Layers** transport authentication data and metadata.

Consider how complex online relationships may become implicit in the IoT. When for instance we operate an Internet-connected domestic appliance, the warranty registration may create a contract with the manufacturer which entitles them to automatically receive information about our usage of the device. When a connected appliance is activated (or new software installed), new and possibly un-noticed relationships can be instantiated between consumers and manufacturers, akin to the "click-wrap" license agreements and terms & conditions of traditional apps. It can be in the interests of manufacturers and other data intermediaries for these relationships to go un-noticed by customers, so that concerned individuals do not object to and thus throttle commercially valuable data flows.

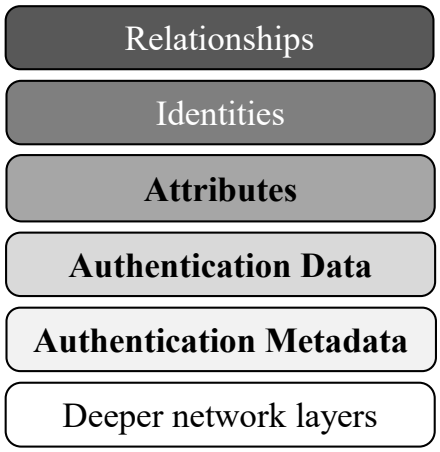

Fig. 1. Schematic digital identity stack

A. Privacy Enhancing Design Patterns in IoT

Conceptualizing IdM according to the digital identity stack helps separate concrete attributes from abstract identities, and leads to design patterns to enhance privacy in the IoT.

One of the most important technical privacy protections is the digital signing of selected data transmitted from devices (a practice which is now spreading, through the FIDO Alliance UAF and U2F protocols, amongst other standards). Conventionally, digital signatures are used to preserve the integrity of messages and identify their originator [21]. However, digital certificates in a public key infrastructure (PKI) can also be used to convey attributes of users rather than their identity [22]. The FIDO attestation certificate [23] is an example of digital signatures being more concerned with the provenance of data and the authentication of attributes, than with the user's identity. By digitally signing transaction data with the FIDO attestation key, a device automatically binds to those data details of the type of device, vouched for by its registered manufacturer. Receivers of signed messages from a FIDO device can be sure of the type of device, and thus make informed decisions about the user in a transaction context, without needing to know who they are exactly.

Where appropriate, digital signatures and attribute certificates can help hide or de-emphasise the personal identity of IoT users, and instead prove particular attribute data to substantiate a transaction. For example, when collecting research data from medical instruments, patient identity is probably irrelevant (or is best suppressed) yet researchers require assurance that patients and their equipment are legitimate and have been correctly enrolled in a study. A digital certificate issued by the study manager and installed in each instrument can specify its approval status without identifying its user [24].

Similarly, in a connected car use case, an entertainment system could digitally sign the listening history before transmitting it to the manufacturer, with a digital certificate that specifies the type of car, without naming the owner. In these cases, the purpose of the digital certificate and signature is to prove the provenance of the data and to associate the data with impersonal attributes of devices, instead of people.

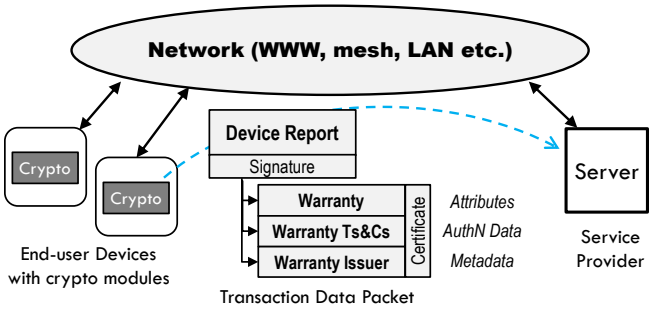

Fig. 2. IoT devices transmitted secure attributes and metadata

Figure 2 illustrates several IoT devices with integrated cryptographic modules to manage private keys and certificates, and digitally sign data sent to other devices or services. One device is shown transmitting a digitally signed report to a server, together with information about the device warranty, the detailed terms & conditions of the warranty, and the name of the warranty issuer. These attributes and metadata can be held in a digital certificate associated with a private key in the device's crypto module. The three layer sub-section of the digital identity stack indicated in Fig 2 shows how the message is confined to attributes, authentication data and metadata, thus minimizing disclosure of personal information about the user. The integrity and provenance of the attribute data may be verified by another device or by a server, first by verifying the digital signature and digital certificate by regular public key cryptography, and then by checking that the warranty details contained in the digital certificate satisfy business rules programmed into the server.

The digital identity stack decouples attributes from personal information, and encourages privacy preserving design patterns:

- designers should use as few low-level attributes as possible when managing IoT devices, to reduce the chances of users being re-identified
- designers should focus on the precise data about devices that are needed to be exchanged with other devices
- there may be more than one source of truth about an attribute; designers should decide which attribute issuer is most appropriate in each context
- use authentication metadata to associate attribute issuers with relevant transactions, so that receivers can more automatically verify the bona fides of users
- metadata can be used to track the history of data flows, showing where each attribute originated and what it was intended to authenticate; secondary re-use of data can then be more constrained, to enhance privacy.

V. Conclusion

We have presented a fresh stack-based model for decomposing and better describing digital identity, with the aim of separating distinct layers of attribute information needed for Identity Management processes. The stack distinguishes high level representations of relationships and identities, from the concrete attributes that make up abstract identities, and further, the low level authentication data and metadata which confer the provenance of attributes. This treatment should help safeguard personal information in the IoT, by focusing designers on precise identity information, stressing which attributes about users and devices are truly relevant in each transaction context, and thus avoiding the over-disclosure of extraneous details which can nevertheless be mined downstream from IoT collection points, to the detriment of privacy. Implementations of the proposed digital identity stack in the IoT can utilize embedded cryptographic modules that are increasingly common in connected devices, and able to digitally sign data, in order to bind attributes and provenance directly to transaction data. Separating attributes and authentication metadata from personal identities should significantly improve privacy in IoT.

References

[1]. Gubbi, J., Buyya, R., Marusic, S. and Palaniswami, M. "Internet of Things (IoT): A vision, architectural elements, and future directions." *Future generation computer systems* 29.7 (2013): 1645-1660.

[2]. Yager, Ronald R., and Jordán Pascual Espada. "New Advances in the Internet of Things." (2017).

[3]. Uden, L and He, W. "How the Internet of Things can help knowledge management: a case study from the automotive domain." *Journal of Knowledge Management* 21.1 (2017): 57-70.

[4]. Tesla Motors Inc. Customer Privacy Policy (undated), https://www.tesla.com/about/legal (accessed 29 August 2017)

[5]. ZUU Online https://zuuonline.sg/business/fintech/data-from-cars-may-be-worth-more-than-the-cars-themselves-in-the-near-future "Data from cars may be worth more than the cars themselves in the near future", July 13, 2017.

[6]. Greenleaf, Graham. "Sheherezade and the 101 data privacy laws: origins, significance and global trajectories." *JL Inf. & Sci.* 23 (2014): 4..

[7]. UK Information Commissioner's Office. Key definitions of the Data Protection Act https://ico.org.uk/for-organisations/guide-to-data-protection/key-definitions (accessed 3 September 2017)

[8]. Office of the Australian Information Commissioner, 2016. Privacy Act. https://www.oaic.gov.au/privacy-law/privacy-act/ (accessed 3 September 2017)

[9]. Organisation for Economic Cooperation and Development. OECD Guidelines on the Protection of Privacy and Transborder Flows of Personal Data. 1980. http://www.oecd.org/sti/ieconomy/oecdguidelinesontheprotectionofprivacyandtransborderflowsofpersonaldata.htm

[10]. Rosner, G.: Privacy and the Internet of Things. O'Reilly, Sebastopol (2016)

[11]. Peppet, S.R.: "Regulating the internet of things: first steps toward managing discrimination, privacy, security and consent." *Texas Law Review*. 93:85 (2014).

[12]. Dewri, R., Annadata, P., Eltarjaman, W. and Thurimella, R. "Inferring Trip Destinations from Driving Habits Data."

12th ACM Workshop on Privacy in the Electronic Society WPES 2013, Berlin, November 4, 2013.
[13]. Yves-Alexandre de Montjoye, César A. Hidalgo, Michel Verleysen and Vincent D. Blondel, "Unique in the Crowd: The privacy bounds of human mobility" Nature Scientific Reports 3, Article number: 1376 (2013)
[14]. Cameron, K.: The laws of identity. Microsoft Corporation (2005).
[15]. FIDO Alliance, 2014. Proposed Standard FIDO UAF Authenticator Metadata Service v1.0, 8 December 2014. https://fidoalliance.org/specs/fido-uaf-v1.0-ps-20141208/fido-uaf-metadata-service-v1.0-ps-20141208.html (accessed 4 September 2017)
[16]. Richer, J. and Johansson L. Vectors of Trust. Internet Engineering Task Force, April 3, 2017 https://www.ietf.org/id/draft-richer-vectors-of-trust-05.txt
[17]. World Wide Web Consortium, Verifiable Claims Working Group, https://www.w3.org/2017/vc/WG/ (accessed 3 September 2017)
[18]. Cavoukian, Ann. "Personal Data Ecosystem (PDE)–A Privacy by Design Approach to an Individual's Pursuit of Radical Control." *Digital Enlightenment Yearbook 2013: The Value of Personal Data* (2013): 89-101.
[19]. Pisa, Michael, and Matt Juden. "Blockchain and Economic Development: Hype vs. Reality." (2017).
[20]. International Telecommunication Union, "Reference Model of Open Systems Interconnection For CCITT Applications", Recommendation X.200 (1988)
[21]. Wilson, Stephen. "Certificates and trust in electronic commerce." *Information Management & computer security* 5.5 (1997): 175-181.
[22]. Wilson, Stephen. "Public key superstructure: it's PKI Jim, but not as we know it!" *Proceedings of the 7th symposium on Identity and trust on the Internet*, pp. 72-88. ACM, (2008).
[23]. FIDO Alliance, 2014. Proposed Standard FIDO UAF Protocol Specification v1.0, 8 December 2014. https://fidoalliance.org/specs/fido-uaf-v1.0-ps-20141208/fido-uaf-protocol-v1.0-ps-20141208.html (accessed 4 September 2017).
[24]. Wilson S. "Anonymity & Pseudonymity in eResearch via smartcards and Public Key Infrastructure". *eResearch Australasia* Sydney, Australia (2009).