Network Working Group
Internet-Draft
Intended status: Informational
Expires: August 22, 2016

S. Farrell
Trinity College Dublin
A. Cooper
Cisco

February 19, 2016

It's Often True: Security's Ignored (IOTSI) - and Privacy too. draft-farrell-iotsi-00

Abstract

Designers of information models for challenged devices connected to the Internet, and most especially for devices that will be carried by people or that will be operating in people's homes, need to not forget that people own the devices and the data, and expect those to work for them, not against them. This draft discusses some security and privacy issues that may be relevant for the IAB's IOTSI workshop on information models for such devices and related services.

Table of Contents

1. Introduction

This is a contribution to the IAB IOTSI workshop. [1] It is not expected to ever become an RFC.

The IETF has recognised the need for strong security mechanisms [RFC3365] to be defined for all IETF protocols. The IETF has further recognised the potential for pervasive monitoring and that work to counter that is needed [RFC7258] and the IAB has produced guidelines for handling privacy in Internet protocols. [RFC6973]. This draft aims to identify some issues with the above that may arise with respect to the information models that are the topic of the IOTSI workshop, but that might othewise get forgotten. Let's start from just a few high-level principles that ought inform designs:

1. Don't forget that the user owns the device and, arguably, the data produced related to that device.

2. Don't forget that the device needs to be updated and that the vendor will end-of-life the device, but the above still needs to be remembered.

3. Don't forget that while we can secure information elements in transit and in storage, that will always be imperfect and information will leak out.

It is worth noting that the IOTSI call for submissions itself did ignore all of these issues.

For each of the above, we'll list a few issues, with some references that may help in further discussion. A full analysis of how each of these ought be reflected in requirements, in specifications of information models and subsequently in data models and protocols is not a goal here, nor is completeness, the goal for now is simply to try ensure that these issues are considered in the IOTSI workshop.

1.1. Ownership and Privacy

1. Regardless of the current legal situation in any particular jurisdiction, it is inevitable that somewhere, sometime, the user will be considered the legal owner of the data emitted by devices relevant to this discussion, sometimes in inventive or disruptive ways. [2] Information models need to not assume that all information elements are "fair game" for all uses by service providers, e.g. no matter how problematic it may be for service providers, explicitly informed consent may be needed to include some data in aggregates. And that might impose a need for what could end up being overly-complex permissions handling for pretty much any element in information models. Managing that without being overwhelmed by complex models may be hard.

2. Don't depend on opaque end-user or legal agreements - users do not, and you are likely to generate terrible publicity if your device or service gets popular. [3] Information models that avoid this error are likely to involve additional entities, such as local controllers that might allow an end-user some control over what their devices are doing.

3. Don't include long-term stable unique identifiers anywhere, and do seriously attempt to avoid all such. You will never know how your information model will be reused or abused. [ishtiaq2010security] While this may seem obvious, it will get forgotten even by people who generally are attempting to be privacy-friendly. In particular using MAC addresses ([I-D.jennings-core-senml] Section 6.1.2) in this way is actively harmful to privacy and in the case of the DHCP protocol, fixing that years later requires a significant specification [I-D.ietf-dhc-anonymity-profile] and implementation effort, and it remains to be seen if such work will get widely deployed or not. It is far better to be highly conservative in the initial stages of work, (where IOTSI is) so that such remedial efforts are not required later.

4. In some circumstances designing systems involving constrained devices involves trade-offs between efficient use of resources and privacy. For example, leveraging hardware identifiers at the application layer may allow for compression or help conserve

bandwidth usage, but may also create additional avenues for attack as compared to using compartmentalized application-layer identifiers. At the point of specifying information models, decisions about how individual systems will navigate these trade-offs should not be taken for granted. Rather, information models should be specified to support privacy-enhancing decisions at the system level, with optional support for less-privacy-enhancing decisions in situations where deployment constraints are expected to warrant such support.

5. Traffic patterns or content, even if "anonymised" can be identifying in unexpected ways, either intrinsically [4] or via correlation. [5] Even the existence/non-existence and timing of application or inftastructure (e.g. DNS, DHCP) traffic can reveal presence or more. Naive information models that don't consider these issues are more likely to result in vulnerable systems.

6. Distinctions between "data" and "meta-data" may not be significant when considering privacy and security - an information or data model or protocol that assumes that e.g. only "data" needs confidentiality is likely broken, as meta-data and traffic patterns may fully breach privacy. There is also a tendency to re-inject data that is carried in ciphertext form into wrappers or headers that are considered meta-data and carried in clear or exposed at too many middleboxes. That anti-pattern is one to be strongly discouraged. [I-D.hardie-privsec-metadata-insertion]

1.2. Life Cycle

1. All devices of any kind will include vulnerabilities. If device software/firmware is not updated, those will eventually be exploited somewhere, sometime. Crawling the network to find those vulnerble devices is a solved problem. [6]

2. The end-user will want the device to continue working and continue getting updated even after all vendors and service providers initially involved have end-of-life'd everything involved. In principle, everything (DNS names, services, roots of trust for software update) needs to be something that can be updated even then. End-of-life is clearly a more complex issue than is typically considered (as shown by the list at [7] which was just a first hit for a search).

1.3. Imperfection

1. We do have ways to protect data in transit and in storage, but we cannot depend on those protecting any information element all of the time. Even with best practices, eventually, some fields from some protocols will leak. All layers need to do as much as possible to provide security and avoid privacy leaks. [8]

2. Even where (structured) data is encrypted, there may still be ways to analyse the traffic to expose the information content. [RFC6562] for example shows that variable bit rate audio with secure RTP can expose audio. And encoded audio is often much more complex than the information considered here.

2. Commercial Considerations

At present, many devices and services are sold and operate in ways that do not take account of the considerations listed here. That is often done for pragmatic and/or commercial reasons, due to the inability to reliably contact devices from the parts of the Internet about which we care, or sometimes in an effort by a vendor or service-provider to achieve "lock-in" so that a user has a hard time mixing and matching the devices and services that the user prefers. And sometimes, users won't have sufficient technical ability to make a device and/or service work for them, even if the vendor or service-provider does expose interfaces allowing for security and privacy friendly deployment.

In this document, the term "service provider" is used consistent with the above, to mean some application service that is not under the control of the device owner or end-user, but rather is controlled by someone else, likely the device vendor or a partner of theirs.

While such cases are a reality and the norm today, and while it is often unclear how to move from there towards a situation where devices and services promote interoperability, the basic information models developed for these devices and services should not preclude a future in which a user can exert independent control over these deployments.

It seems likely from the above that information models will need to include some conception of the device owner as a first-class, but hopefully pseudononymous, entity and not be solely limited to consideration of characteristics of devices and services.

3. Security Considerations

Yes, there are. Are you shocked?

4. Privacy Considerations

Yes, there are. Aren't you shocked yet? :-)

5. IANA Considerations

This document makes no requests for IANA action. This section would be removed except it won't be as we're not aiming for publication as an RFC.

6. Acknowledgements

TBD - your name here for comments or beer!

7. References

7.1. Informative References

[I-D.hardie-privsec-metadata-insertion]
Hardie, T., "Design considerations for Metadata Insertion", draft-hardie-privsec-metadata-insertion-00 (work in progress), October 2015.

[I-D.ietf-dhc-anonymity-profile]
Huitema, C., Mrugalski, T., and S. Krishnan, "Anonymity profile for DHCP clients", draft-ietf-dhc-anonymity-profile-07 (work in progress), February 2016.

[I-D.jennings-core-senml]
Jennings, C., Shelby, Z., Arkko, J., and A. Keranen, "Media Types for Sensor Markup Language (SENML)", draft-jennings-core-senml-04 (work in progress), January 2016.

[RFC3365] Schiller, J., "Strong Security Requirements for Internet Engineering Task Force Standard Protocols", BCP 61, RFC 3365, DOI 10.17487/RFC3365, August 2002, <http://www.rfc-editor.org/info/rfc3365>.

[RFC6562] Perkins, C. and JM. Valin, "Guidelines for the Use of Variable Bit Rate Audio with Secure RTP", RFC 6562, DOI 10.17487/RFC6562, March 2012, <http://www.rfc-editor.org/info/rfc6562>.

[RFC6973] Cooper, A., Tschofenig, H., Aboba, B., Peterson, J., Morris, J., Hansen, M., and R. Smith, "Privacy Considerations for Internet Protocols", RFC 6973, DOI 10.17487/RFC6973, July 2013, <http://www.rfc-editor.org/info/rfc6973>.

[RFC7258] Farrell, S. and H. Tschofenig, "Pervasive Monitoring Is an Attack", BCP 188, RFC 7258, DOI 10.17487/RFC7258, May 2014, <http://www.rfc-editor.org/info/rfc7258>.

[ishtiaq2010security] Ishtiaq Roufa, R., Mustafaa, H., Travis Taylora, S., Xua, W., Gruteserb, M., Trappeb, W., and I. Seskarb, "Security and privacy vulnerabilities of in-car wireless networks: a tire pressure monitoring system case study", 2010.

7.2. URIs

[1] https://www.iab.org/activities/workshops/iotsi/

[2] https://www.economist.com/blogs/schumpeter/2014/06/who-owns-your-personal-data

[3] http://techcrunch.com/2015/02/08/telescreen/

[4] http://www.economist.com/blogs/schumpeter/2014/06/who-owns-your-personal-data

[5] https://en.wikipedia.org/wiki/Netflix_prize#Privacy_concerns

[6] https://www.shodan.io/

[7] https://www1.good.com/support/end-of-life-notices.html

[8] https://en.wikipedia.org/wiki/AOL_search_data_leak

Authors' Addresses

Stephen Farrell
Trinity College Dublin
Dublin 2
Ireland

Phone: +353-1-896-2354
Email: stephen.farrell@cs.tcd.ie

Alissa Cooper
Cisco
707 Tasman Drive
Milpitas, CA 95035
USA

Email: alcoop@cisco.com