



HAL
open science

The Responsibility of Open Standards in the Era of Surveillance

Harry Halpin

► **To cite this version:**

Harry Halpin. The Responsibility of Open Standards in the Era of Surveillance. Hot Topics in Privacy Enchancing Technologies, Jul 2016, Darmstadt, Germany. hal-01426848

HAL Id: hal-01426848

<https://inria.hal.science/hal-01426848>

Submitted on 5 Jan 2017

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

The Responsibility of Open Standards in the Era of Surveillance

Harry Halpin
W3C/INRIA
harry.halpin@inria.fr

1. THE ROLE OF OPEN STANDARDS?

The core infrastructure of the Internet is defined by interoperability between code-bases: The ‘rough consensus and running code’ of open standards at the Internet Engineering Task Force (IETF)¹ and World Wide Web Consortium (W3C).² However, there are a number of powerful critiques of open standards. First, there is a widespread failure of many core standards in terms of security and privacy, and even concerns of subversion. There is an even more substantial critique that standards are simply moving too slowly in the face of rapid innovation. However, we’ll argue that engagement with open standards is the best way for privacy-enhancing technologies to gain widespread adoption.

What are open standards? Open standards are defined by IETF/W3C Working Groups, who provide both the formal structure and patent agreements necessary for engagement with many parts of industry, as well as ideally long-term maintenance of the specified protocol. In contrast, Tor is at present not a standard. Only if there were multiple Tor code-bases or forks that needed to communicate would a standard be needed. Officially, at both the W3C and IETF there must be two interoperable implementations for any draft specification to officially become a standard.

Open standards are defined as “open” in terms of participation, in contrast to “closed” standards bodies such as the ITU or ISO where participation requires government status. Also, open standards bodies also publish their standards free of charge, encouraging anyone to implement. While open standards are typically required by commercial actors for anti-trust reasons, open processes also tend to be good practice from a security perspective, as the review of multiple experts typically discovers security flaws. This is in contrast to the rather mysterious process of standardization at many national-level standards bodies, where the lack of transparency can lead to subversion. The IETF has no formal membership (so that one “joins” by participating on a mailing list). The W3C has a formal membership and invited expert status for the general public. Typically, academics and activists can participate, although in practice participation is limited due to time, travel, and lack of academic incentives.

Rather than code licensing, open standards deal with patents. A patent holder can still claim patent infringement even if an idea is embodied in free software. W3C standards are explicitly licensed by W3C members under a royalty-free license.³

¹<http://www.ietf.org/>

²<http://www.w3.org/>

³<https://www.w3.org/Consortium/Patent-Policy->

In contrast, the IETF “Note Well” policy simply requires disclosure of known patents.⁴ The much stronger W3C policy essentially creates a kind of “patent war-chest” composed of all W3C standards, from XML to HTML5. This patent war-chest is then enforced by a “balance of terror” so that any member that makes a patent claim on a W3C standard triggers their loss of royalty-free licensing for all W3C standards. Given the historical role that patents have had in holding back the development of cryptography, ranging from the RSA to Schnorr to Certicom patents, open standards should be a pre-requisite for cryptographic protocols.

The last benefit of open standards is long-term governance. Typically, programs are made either by for-profit companies that may eventually collapse or loose networks of programmers where crucial decisions are made by a few informal leaders. Despite all the churn of companies and open-source projects, standards bodies like the IETF and W3C have so far survived and preserved the Net and Web as co-operatively governed platforms defined by open standards. So open standards are one path that those who are unaffiliated with a corporation can use to influence the entire infrastructure of the Net.

2. FAILURES OF STANDARDIZATION

Yet there is disenchantment with open standards from developers of privacy-enhancing technologies. In terms of governance, there is a widespread belief that existing standards could be subverted. The NSA clearly subverted national-level standards such as Dual EC DBRG of NIST, and the military-industrial complex has long been close to the IETF, as witnessed by the scandal caused by the co-chair of the IRTF Crypto Forum Research Group (CFRG) being an NSA employee (who eventually quietly stepped down, although he was not removed).⁵ As shown by the long process to standardize a replacement elliptic curve for P-256 in the CFRG, even an open and informal process may not be considered fast and fair by some developers.

In terms of quality, current standards with mass deployment are considered often poorly-designed, with privacy and security ‘bolted on’ as an afterthought. Two well-known cases are the Battery API leaking private information[2] and TLS having widely reported security vulnerabilities[1]. Worse, standards may even spread technologies that are insecure and privacy-invasive: The W3C is standardizing access to DRM in Encrypted Media Extensions against the

[20040205/](http://www.w3.org/20040205/)

⁴<http://www.rfc-editor.org/rfc/rfc3979.txt>

⁵<http://www.ietf.org/mail-archive/web/cfrg/current/msg03554.html>

recommendations of many in the security community.⁶ The IETF/W3C WebRTC reveals the IP address of its user even if a VPN is being used: The standards process ignored the needs of activists, like in Iran, that were using VPNs to obscure their identity.

The open nature of standards is ultimately the best defense against subversion. To a large extent these failures of the standardization process result from a failure of participation, with errors being noticed only after deployment of the standard. Without the right experts at the table from the start, flaws will not be detected until after standardization. The W3C and IETF are welcoming to more security and privacy experts joining Working Groups. After all, the IETF proclaimed that “pervasive monitoring is an attack” after the Snowden revelations.⁷ There has been progress as well, as new standards like TLS 1.3 and Web Authentication are aggressively improving security. Yet privacy reviews are still for the most part voluntary and there is a distinct lack of awareness of data minimization.

A more fundamental critique of standards is that open standards freeze innovation. Marlinspike argued that standardized federation like XMPP and PGP should be abandoned in favor of more rapid innovation.⁸ Rapidly pushing out updates (for not only user experience but also protecting metadata) seems to go against the slow process of standardization. While this critique makes sense in terms of security updates, it only makes sense for the platform as a whole if one believes that the primary driver of innovation is the original software development team. As shown by the rapid proliferation of app-stores and open APIs, closed platforms tend to open up in order to capitalize on user-driven open innovation. Only core functionality should be standardized in order to enable open “permissionless” innovation to happen both above and below the standard. Standards only “freeze” development if extensibility points are not well-defined and the standards body fails to continue to upgrade the standard in the face of innovation. Today’s growth of closed silos on top of open standards is a testament to both sides of this principle.

Although the use of the address-book as the “portable” social network in mobile devices is a (privacy-invasive) starting point, it is far from the only layer that is in need of standardization. Is there a way to boot-strap new applications without handing the entire social graph to a third-party server? Can users save and move media and messages from one (encrypted) encrypted messaging app to another? These tasks require open standards. As the Internet is currently divided into mutually incompatible silos whose network effects make it difficult for start-ups like Open Whisper Systems to gain enough users to take control of the market, open standards are key for letting users switch applications easily and even communicate to existing silos.

3. THE FUTURE OF STANDARDS

One should not be in despair of the current dismal state of standards in domains like secure messaging: The slow progress of XMPP (OTR) and SMTP (PGP) is due to attempts to add security to protocols that were never built

with security in mind. What is needed is new standards with correct data minimization properties and security guarantees from the beginning. Yet what new standards should be pursued? Despite all the talk of innovation, there are only a small number of tasks that compose everyday Net use: Messaging, file-sharing, calendaring, and the like. Could each of these have a new privacy-enhanced protocol? In order to chose new candidates for standardization a number of simple criteria should be applied to new standards efforts: 1) number of open-source codebases as well as closed-source products pursuing similar goals 2) number of competing protocols in the space used by aforementioned projects/products 3) number of users of these protocols, with a bias to users in at-risk areas 4) relative maturity of the protocols themselves.

A new generation of protocols could be adopted if either privacy is ultimately valued as a fundamental right (via popular or regulatory pressure) or a “privacy market” emerges. Already the IRTF Human Rights Protocol Considerations Research Group⁹ is investigating the connections between privacy-enhanced protocols and the Apple vs. FBI case has demonstrated market interest in encryption and privacy. Concretely, there are a number of candidate protocols that would be required for end-to-end encrypted messaging across various platforms, including key discovery[3] and end-to-end encrypted messaging.¹⁰ We should begin a new post-PGP standardization effort that brings privacy to everyone, not just the cryptographic elite.

Due to its foundation as an academic network and historical restrictions on cryptography, the Internet was built without strong cryptography and privacy. As put by Vint Cerf, “I worked with the National Security Agency on the design of a secured version of the internet but we used classified security technology at the time and I couldn’t share that with my colleagues. If I could start over again I would have introduced a lot more strong authentication and cryptography into the system.”¹¹ After Snowden, the historical responsibility of open standards bodies like the IETF and the W3C is to upgrade the common infrastructure. However, they will not be able to do so without a new generation of privacy academics and activists guiding these standardization efforts. This historical task requires going beyond deploying privacy-enhancing technology only in particular companies, start-ups, or code-bases, but to recognize our responsibility to build privacy into the open standards themselves.

4. REFERENCES

- [1] K. Bhargavan, A. D. Lavaud, C. Fournet, A. Pironti, and P. Y. Strub. Triple handshakes and cookie cutters: Breaking and fixing authentication over tls. In *Security and Privacy (SP), 2014 IEEE Symposium on*, pages 98–113. IEEE, 2014.
- [2] C. Diaz, L. Olejnik, G. Acar, and C. Castelluccia. The leaking battery: A privacy analysis of the html5 battery status api. *Lecture Notes in Computer Science*, 2015.
- [3] M. S. Melara, A. Blankstein, J. Bonneau, E. W. Felten, and M. J. Freedman. Coniks: bringing key transparency to end users. In *24th USENIX Security Symposium (USENIX Security 15)*, pages 383–398, 2015.

⁶<http://boingboing.net/2016/03/29/security-researchers-help-eff.html>

⁷<http://www.rfc-editor.org/rfc/rfc7258.txt>

⁸<https://whispersystems.org/blog/the-ecosystem-is-moving/>

⁹<https://irtf.org/hrpc>

¹⁰<https://github.com/whispersystems/>

¹¹<http://now.avg.com/an-insiders-look-at-the-history-of-cybersecurity/>