

Reviewing for Privacy in Internet and Web Standard-Setting

Nick Doty

UC Berkeley, School of Information
Berkeley, USA

Email: npdoty@ischool.berkeley.edu

Abstract—The functionality of the Internet and the World Wide Web is determined in large part by the standards that allow for interoperable implementations; as a result, the privacy of our online interactions depends on the work done within standard-setting organizations. But how do the organizational structure and processes of these multistakeholder groups affect the engineering of values such as privacy? This paper reviews the history of considerations for security and privacy in Internet and Web standard-setting; the impact of Snowden surveillance revelations and reactions to them; and some trends in how we review for privacy in Internet and Web standards.

Keywords—Internet, Privacy, Standards organizations, World Wide Web.

I. INTRODUCTION

The functionality of the Internet and the World Wide Web is determined in large part by the standards that allow for interoperable implementations; as a result, the privacy of our online interactions depends on the work done within standard-setting organizations. While privacy impact assessments present a systematic model for privacy reviews of large-scale software systems, reviewing Internet standards provides a different set of challenges. Participation in consensus standard-setting is voluntary and in most cases work is bottom-up, unlike the top-down organization within large firms. And Internet protocols are layered and generative: designed to enable a variety of applications, they are resistant to static analysis. In prior work we have analyzed these multistakeholder groups as “boundary organizations” which can provoke innovative responses [1]; this organizational structure provides a different approach to implementing values such as privacy that is not yet fully understood.

In Section II, I describe the methods, data and scope of the paper. In Section III, I first detail how process, tools and organizational structure affected security considerations at IETF and the relation to privacy in Internet standards. Next, in considering privacy-specific standards for the Web, I show how privacy is conceived for that application and how privacy reviews are currently done. In Section IV, I describe the different reactions in the standard-setting community to the Snowden revelations and their effect on reviewing for privacy. Based on that history, in Section V I identify three significant trends — systematization, integration and leadership — in reviewing for privacy in Internet and Web standard-setting.

II. METHODS

A. Data Sources

This paper documents the historical practice and current trends of reviewing for privacy (and, related, security) in standard-setting through three data sources. First, the Internet and Web standards themselves provide a corpus of text documents for automated text analysis, which can indicate and confirm trends quantitatively; related, activity on publicly-archived mailing lists is collected. Second, mainstream news articles, meeting reports and key standards documents are used to detail the timeline and character of responses to Snowden revelations. Finally, semi-structured interviews with Internet engineering experts and participants in IETF activities are used to provide an internal perspective on processes.

B. Scope

This paper does not delimit or assume a single definition of privacy. My research is focused privacy as encompassing freedom from intrusion and control over information about oneself. But because privacy is an essentially-contested concept [2][3], I have sought not to prime interviewees or assume that all software engineers or standard-setting participants have a common view of privacy. How privacy is conceived by those individuals may in fact have a substantive effect on the privacy outcomes for Internet and Web users.

Collection and analysis of published standards, documents, mailing list conversations and participant interviews are focused on two technical standard-setting bodies, the Internet Engineering Task Force (IETF) and World Wide Web Consortium (W3C), consortia that use voluntary, consensus-based models and function as prominent venues for work on Internet and Web protocols.

Because of the voluntary and informal nature of these standard-setting bodies, there are not sharp boundaries delimiting where standards of Internet-related technologies are discussed. For example, many Web standards (notably including HTML5) are substantially developed within the Web Hypertext Application Technology Working Group (WHATWG), an organization formed by a group of browser vendors [4]. That work can be conducted separately from, or in concert with, W3C discussions, through an “uneasy collaborative alliance” [5]. OASIS (formerly SGML Open) has developed standards for applications of Extensible Markup Language (XML) to healthcare, security, web services and business processes [6]. Recently, OASIS Technical Committees have developed stan-

dards for organizational methodology for handling privacy [7] and are discussing privacy-by-design process standards [8].

The list of standard-setting venues is too long to describe in the space of this paper.¹ Supported by my access to IETF and W3C activities, this analysis will primarily follow reviewing for privacy at IETF and W3C. As described in *Future Work*, these multistakeholder groups may provide insight into similar organizational settings; whether the same patterns apply in different kinds of standard-setting, or where they differ, would be a useful point of comparison.

III. HISTORY OF PRIVACY AND SECURITY REVIEWS

For this paper, first I consider the history of privacy reviews and, related, security considerations at IETF and W3C through analysis of documents and interviews with participants.

A. Internet Engineering Task Force

IETF has an extensive history of published documents to analyze, since 1969. While privacy may not have been an explicit topic in those early standards, Sandra Braman has suggested that the basic value of privacy is apparent in many early IETF standards publications — called Requests for Comment (RFCs) — often described in terms of communications security [12].

Security considerations provide a useful precursor for studying reviewing for privacy for multiple reasons. Security — like privacy, accessibility, internationalization, performance or other values — is a cross-cutting, cross-functional or “horizontal” concern. For various conceptions of privacy, security properties (confidentiality or integrity, in particular) are prerequisites; for “layered” technologies, security properties at a lower layer may determine whether privacy can be enabled for a particular application. And while security is related to privacy, it’s also an area with which there is more experience and methodology we might adapt to support of privacy [13].

To that end, in the next section I look at the history of Security Considerations sections in IETF RFCs and how they’ve changed over time; in the following section I briefly describe more recent IETF work that explicitly considers privacy.

1) *Security Considerations*: After a process requirement was added such that all RFCs were mandated to have a Security Considerations section, we see a dramatic increase (with almost complete compliance) to at least mention security; see Fig. 1. However, mentions of “privacy” don’t see the same marked increase. While mention of “security” reaches approximately 100% after 1990, mentions of “privacy” seem to level off at a fifth of all the documents published. This quantitative measure supports the notion that mandates in a standard-setting process make some difference in the resulting published documents.

¹For an example of the variety: the Kantara Initiative (mentioned later) hosts working groups to discuss Internet identity management technologies [9]; Ecma is known for work on the JavaScript language widely used on the Web [10]. More formal standard-setting organizations include the International Organization for Standardization (ISO) which counts national standard-setting bodies as its members and the International Telecommunications Union (ITU), a United Nations specialized agency which coordinates spectrum usage and develops telecommunications standards [11].




Fig. 1. Percentage of published RFCs with search terms, by year. The lower graph shows the number of RFCs published by the RFC Editor each year from 1969 to 2014. The labeled colored lines in the upper graph show the percentage of the documents published each year that make at least a single mention of the particular term.

The formal requirement for having a security considerations section is present first in RFC 1543, published in 1993 and titled simply “Instructions to RFC Authors”. That RFC is essentially a style guide, describing the format of RFCs for publication, and continues to be updated in that way, most recently in September 2014 by the RFC Editor [14]. Being purely about style, however, it provides neither requirements nor guidance on the contents of a security considerations section. It’s common knowledge in IETF circles that Security Considerations sections in the 1990s were typically insufficient. For example, in the abstract of a 2003 RFC providing detailed guidance on Security Considerations sections [15]:

All RFCs are required to have a Security Considerations section. Historically, such sections have been relatively weak.

In a selection of RFCs from 1996, Rabkin [16] found few substantive Security Considerations sections and found that all were brief.

An automated text analysis of RFCs confirms and extends those findings; see Fig. 2.² Security Considerations sections are absent prior to 1990 and tend to be very short in RFCs published in the 1990s. In the past 15 years, many Security Considerations sections are longer and represent a larger fraction of the length of published documents, although many minimal-length Security Considerations sections remain. The length of Security Considerations text is no guarantee of good security, but this measure does indicate increased attention and thought, while also noting that some documents published today have little to say about security.

The basic level of compliance is actually enforced technically. Templates for writing Internet-Drafts include a Security Considerations section and submissions tools check for the

²Numbers in this graph should be considered estimates, as the lack of consistent formatting of RFCs over the past 40 years creates error in the automated detection of sections. Code is available; review and improvements would be welcome: github.com/npdoty/rfc-analysis




Fig. 2. Length of Security Considerations sections in RFCs between 1969 and 2014. Generated by code parsing the plain-text representation of RFCs through 2014 and identifying sections automatically based on text formatting; length is determined by number of fixed-width lines. Dots colored more orange represent documents where the Security Considerations section represents a larger fraction of the total lines of the document; blue represents a smaller fraction.

presence of such a section [17]. The substance of those sections is supported by assigned reviews conducted by volunteers who participate in a Security Directorate [18]. During interviews, one participant at the time described an initial motivation that gathered participants for the Security Directorate as the free lunches provided, but “once it was institutionalized and organized, [...] there was enough momentum to keep it going.”

Importantly, leadership is also credited with enforcing the substance of Security Considerations sections. As one current IETF participant put it: “Now everyone [thinks about security]. Not everyone does, but as soon as you don’t, you get called out.” Credited with that calling out are the Area Directors, and in particular the two Security Area Directors. Because every potential RFC is reviewed by the IESG (Internet Engineering Steering Group) before publication, documents may be rejected or subject to revisions if the security (or indeed, privacy) considerations are found lacking. Such reviews can be detailed and approval difficult to obtain in some cases; “the security area directors are like a force to be reckoned with at this point.”

2) *Expanding to Privacy*: Where RFC 3552 provided guidance on writing security considerations in IETF documents, RFC 6973 (drafted over several years and published in July 2013) has attempted to do the same for privacy in Internet protocols: listing specific kinds of threats, mitigations for those threats and a checklist of questions for identifying and addressing privacy issues [19].

As seen below for W3C, there are also IETF standards that take privacy as an explicit substantive aim. In particular, the `geopriv` Working Group (recently concluded) developed standards for transmission of location information that considered user participation and expression of policies — privacy not just in the sense of communications security — as key deliverables [20]. Over several years, the Working Group developed requirements, threat analyses, file formats, architectural plans around a basic model of communicating geolocation information in a standard format with the requirement to




Fig. 3. Percentage of Technical Reports (TRs) published by W3C with search terms, by year. The lower graph shows the number of TRs published each year from 1996 to 2014. The labeled colored lines in the upper graph show the percentage of the documents published each year that make at least a single mention of the particular term. Note: TRs include in progress Working Drafts for recent years; the lower graph doesn’t indicate dramatically increased output.

communicate user-controlled policies for how that location data could be used. For some applications and implementations (for use on the Web, in particular), this model was controversial and rejected by those who found it too complex [21].

That this more explicit privacy focus is found at IETF in a more application-focused area is consistent with the privacy-specific standards described below at W3C; as the Web is a popular application built on top of the Internet.

B. Privacy in World Wide Web Consortium Standards

W3C also publishes technical standards in publicly-accessible documents we can analyze. Without specific process requirements (like the mandated Security Considerations at IETF), the fractions of documents mentioning privacy and security terms seem to be relatively stable (around 20% and 35%, respectively); see Fig. 3, in comparison to Fig. 1.

1) *Privacy-specific Standards*: At W3C, some significant standards efforts have specifically addressed privacy. That is to say, not only do those standards include discussion of the privacy implications or considerations in the development of a new technology, but in fact the standard is itself intended to address privacy on the Web.

a) *Platform for Privacy Preferences Project*: The Platform for Privacy Preferences Project (P3P) was a multi-year effort to improve awareness of Web privacy practices through machine-readable descriptions of Web site privacy policies. Other attempts to address the problem of inadequate privacy notices proliferate today, as can be seen by the regular interest in privacy icons projects³ or US government supported multistakeholder efforts to establish standards for

³Solon Barocas maintains the most exhaustive list of projects in this area: 57 at the time of this writing [22]. The Open Notice (<http://opennotice.org>) group coordinates multiple projects in the same area, and itself is developing standards within the Kantara Initiative Information Sharing Working Group.

short-form transparency notices.⁴ P3P defines a descriptive, extensible XML language for communicating site privacy practices, expanding on a concept previously used for describing content appropriateness ratings for different age groups.⁵ The design of P3P was explicitly layered to provide a neutral, descriptive language for practices and to provide flexibility for implementers or users to draw their own conclusions about the importance of various practices and the differences in their privacy preferences [24]. As we have noted previously, this embodies the frequently-cited technical design principle of “mechanism, not policy” [25]. While machine-readable policies were implemented by some Web sites (especially compact policies in response to an Internet Explorer cookie-blocking policy), P3P never saw widespread use by sites or by browsers, attributed to its complexity or to lack of incentives [26].

b) Do Not Track: The development of Do Not Track (DNT) standards is a multi-year effort to provide a user choice mechanism for tracking of online behavior. After the idea was endorsed by a Federal Trade Commission staff report in late 2010 [27], preliminary deployments of a Do Not Track header flag began in browsers and efforts to standardize began during 2011. DNT is designed to allow users of the Web to indicate simply a preference for or against tracking via their browser (or “user agent”) software and have that preference communicated to all online services. Like P3P before it, DNT doesn’t itself enforce any privacy properties or automatically limit tracking activity, but relies on a cooperation between user agents and servers that choose to comply with those user preferences.

More recently, the technical DNT mechanism has been updated to allow servers flexibility in indicating how they comply with a user’s preference. The Electronic Frontier Foundation (EFF) has developed a Privacy Badger tool that blocks cookies except where it sees the verbatim presence of a proposed DNT policy [28]. The Digital Advertising Alliance (DAA) has indicated its intention to convene its own process for developing a DNT (or similar browser-based opt-out tool) policy [29]. Edited in part by this author, the W3C Tracking Protection Working Group (TPWG) is completing work on a compliance policy to which sites may adhere [30]. This divergence of policies is characteristic of a lack of standardization, or, more specifically, of narrowing standardization such that implementations can increasingly vary.

2) Privacy Reviews at W3C: Privacy issues arise in the specification of various Web APIs and protocols that aren’t privacy-specific. While W3C process doesn’t have formal requirements for the presence of security or privacy considerations in a particular section of published documents, there are steps in the process and formal or informal organizations for conducting broad reviews. A “Last Call” or “wide review” phase (present in IETF and W3C standardization) requests broader review of specifications outside just the authors or Working Group developing the specification. Liaisons are detailed in the charters of individual Working Groups to

pre-define other groups to ask for review comments, often including privacy, security, accessibility or internationalization. The W3C Director has discretion over whether a specification will proceed along the standards track and has the option to require further work to address identified concerns.

On an informal or consulting basis, the Privacy Interest Group (PING) provides advice or reviews of privacy issues with various specifications, typically when a Working Group has specifically solicited them to do so. PING (organized in part by this author) is made up of volunteers from academic, civil society or industry organizations with a particular interest in Web privacy. Those volunteers also collaborate on documents and processes for improving privacy reviews (see *Systematization*, below). Other W3C groups also provide feedback on privacy and security issues: for example, when members of the Technical Architecture Group (TAG), a small group of experts providing oversight, have a particular privacy or security interest, they may comment on those aspects of technical architecture. The Web Application Security Working Group and Web Security Interest Group often publish relevant documents or have interested individuals who provide comments on various Web APIs in progress (see *Integrating Privacy and Security* below).

Privacy is not the only cross-cutting concern that has led to reviews in W3C standard-setting and some of those areas are more formally developed. The Internationalization Working Group⁶ provides advice and reviews to groups inside and outside of W3C on the usability of technologies in different languages. The Web Accessibility Initiative⁷ both develops its own standards and provides reviews of other W3C work related to use of Web technologies by people with disabilities.

IV. REACTIONS TO SNOWDEN

Revelations of widespread government surveillance of electronic communication have profoundly affected the Internet standard-setting community. Reactions over the past 18 months have significantly included: individual and organizational statements; the formation of new groups and collaborations; and direct responses to surveillance in both standards and code. A full description and analysis of those events would fill many papers;⁸ here I provide a brief summary, focused on the impacts for privacy and security reviews of standards in the future. As described above, reviewing for privacy in standards was a practice before any publications regarding Edward Snowden. However, this exogenous event has inspired concrete responses and changed the practices of standard-setting organizations.

The first Greenwald articles based on Snowden sources were published in June 2013 (at the same time as the Privacy Law Scholars Conference in Berkeley), providing details on Section 215 telephony metadata collection and *Prism* access to servers at large tech companies [31][32]. But more relevant to those who work on securing Internet connections themselves

⁴The National Telecommunications & Information Administration (NTIA) maintains a web page on the Privacy Multistakeholder Process: Mobile Application Transparency [23], including a 2013 draft code of conduct.

⁵The Platform for Internet Content Selection (PICS) was published in the late 1990s and later superseded by the Protocol for Web Description Resources (POWDER), developed between 2007 and 2009.

⁶w3.org/International/core/

⁷w3.org/WAI

⁸While no doubt related to standard-setting, this paper does not describe reactions in the area of Internet governance, which would include: the Montevideo I* statement, the NetMundial Initiative and IANA transition proposals.

were XKeyscore, revealed in July [33], and Bullrun, revealed in September [34], which provided surprising evidence of National Security Agency (NSA) capabilities and practices for surveilling Internet activity, including encrypted traffic. Even more specific to the standard-setting community, the same September article provided confirmation that the NSA covertly introduced security vulnerabilities in the development of a technical standard for encryption at the National Institute of Standards and Technology (NIST).

Some responses in the professional community of technical standard-setting were emotional in nature. The full text of the seven-page “A Simple Statement”, an Internet-Draft from a prominent individual contributor to IETF standards, is included below [35]:

we had a good thing
you messed it up
for everyone
we trusted you
we were naive
never again

The IETF’s November 2013 meeting contributed to broader, organizational statements. In a plenary session with several hundred attendees [36], Russ Housley, Chair of the Internet Architecture Board (IAB),⁹ asked for support for, or opposition to, the following statement:

Pervasive surveillance is an attack, and the IETF needs to adjust our threat model to consider it when developing standards track specifications.

The room provided a strong “hum” in favor of the statement, and silence in opposition to it [37]. That consensus was later reflected in a more thorough document [38], detailing the nature of the attack and the process for IETF mitigations. In particular, RFC 7258 notes that considerations for the threat of pervasive monitoring must be present in the technical design of both new and existing protocols, but that a separate “pervasive monitoring considerations” section isn’t necessary.

In addition to individual and consortium-level statements, interested individuals found and formed groups in reaction to the surveillance revelations. The XKeyscore news was published in July [33] during the IETF meeting in Berlin; an informal group met there, which spawned the *perpass* mailing list (the most basic formal organization in IETF and many engineering groups), a BoF (“birds of a feather” meeting for collaboration prior to formal standard-setting activity) at the IETF meeting in November,¹⁰ and a workshop (organized by IAB and W3C) on strengthening the Internet against pervasive monitoring the following February [40].

Mailing list statistics can provide a coarse perspective on levels of activity within a community. The *perpass* list (in orange) shows a dramatic spike in late 2013, relative to other lists. The timeline indicates that this self-organized conversation began after Snowden revelations and became most active after XKeyscore and Bullrun announcements




Fig. 4. Moving average of activity on IETF and W3C mailing lists focused on privacy and security issues.

that described specific subversion of Internet infrastructure. Topics include brainstorming and critiquing of proposals for increased use of encryption; the possible creation of new Working Groups to develop new security standards; discussions of threat models and responses (including what would become RFC 7258); and discussion of processes for conducting privacy reviews.

Perhaps the most concrete responses to surveillance revelations have been pro-active shifts towards encrypting online communication. That trend includes individual companies encrypting their internal communications: for example, Google encrypted data center links [41] in response to revelations of the NSA Muscular program [42]. At the level of Internet and Web protocols, there has been a concerted shift to make more Web browsing encrypted. For W3C, the TAG has outlined a finding of steps towards moving the Web to HTTPS browsing [43]. The IAB has published a statement on confidentiality, encouraging encryption at all available levels for new protocols (sometimes summarized as “no new cleartext”) [44]. While not included as an interoperability requirement, the new HTTP/2 protocol has seen announcements from browser vendors that it will be used only over authenticated, encrypted connections.¹¹

The prevalence of network-level attacks as described by Snowden may also have spurred privacy discussions about the use of various Web APIs. Some have advocated for restricting privacy-sensitive features in the browser (for example, APIs for accessing location, camera or other sensors) to only those pages loaded over secure connections [46]. Those proposals have been debated within several W3C Working Groups. Motivations include both securing connections for those sensitive actions and providing an extra incentive for site developers to deploy secure connections.

V. TRENDS

A. Systematization

As seen from the example of security considerations at IETF, experience and guidance apparently improve the sub-

⁹A small advisory group to the IETF, selected from and by IETF participants.

¹⁰The *perpass* BoF meeting agenda gives an overview of the topics discussed there [39].

¹¹See the HTTP/2 FAQ, for a brief distillation of a long discussion [45].

stance and significance of these reviews. We expect — and are at work on — the same in the area of privacy.

The IAB has attempted to introduce RFC 6973 through educational tutorials for chairs of IETF Working Groups and other interested participants [47]. For W3C, individuals have created similar checklists for spotting potential privacy issues in new Web APIs [48] and guidance for mitigating common privacy issues, such as fingerprinting (edited by this author [49]). But discussion also includes more process-oriented proposals, including a Privacy Specification Assessment [50]. More like privacy impact assessments as seen in government and corporate processes (for example, see [51]), this would describe roles, workflows and timing for identifying and addressing privacy issues throughout a longer lifecycle of designing and implementing a specification.

While systematization continues, it's not without difficulties. Apparently because W3C standards are developed in such a distributed fashion, there's been more success getting informal advice and review from the Privacy Interest Group than in developing a formal process for privacy throughout the specification's design. While IETF has a Security Directorate, a Privacy Directorate was dropped for lack of activity.¹²

B. Integrating Privacy and Security

Perhaps as a response to the challenge of systematic volunteer review of specifications in a non-hierarchical standard-setting environment, privacy and security reviews are increasingly being done at once, rather than by separate groups or individuals. As a logistical matter it may be easier to recruit one individual to review the security and privacy considerations of a particular document, rather than requiring separate people to get up to speed on a technology in order to provide useful comments. Furthermore, since many privacy properties require the security of underlying protocols and much of the thought process for identifying privacy and security issues is similar, there may be significant substantive overlap.

At W3C, this integration is proposed by potentially sharing the review workload between the Privacy Interest Group and the corresponding Web Security Interest Group [53]. The IAB has recently replaced the previous Security Program and Privacy Program¹³ with a single group discussing confidentiality, resiliency and trust. Addressing the threat of pervasive monitoring at IETF has led to proposals for conducting both retrospective and prospective reviews.¹⁴ The focusing event of revelations of mass surveillance of Internet activity may inspire reviewing for privacy in a way that prioritizes that particular threat.

C. Leadership

Technical standard-setting for the Internet and the World Wide Web is famously (or infamously) *not* hierarchical in the form of traditional firm organization. Quoting David Clark, a basic tenet of IETF process has been [56]:

¹²See thread on [privacydir] Closing ML [52]. In the mailing list activity graph above, the red line shows minimal usage of the short-lived privacydir list.

¹³The IAB Privacy Program had coordinated eight privacy reviews [54].

¹⁴Discussed on the perpass mailing list [55], and subsequently at the IETF 89 meeting.

We reject kings, presidents and voting. We believe in rough consensus and running code.

Despite this disavowal of formal power mechanisms, leadership plays an important role in the development of technical standards and is significant in the development of privacy and security considerations. Initial interviews with IETF participants identify the seniority of Area Directors and the process of IESG approval as essential to security and more recently privacy considerations in Internet standards. Statements from quasi-leadership organizations have been prominent in responding to Snowden revelations.

However, individuals and organizations have pushed back against leadership, both within standard-setting and against the governance role of standards. Vendors of “middleboxes” (a generic term for proxies, caches, Internet Service Providers and others that sit in the “middle” of network connections) have objected to proposals for end-to-end encryption of Internet communication, despite the apparent consensus for increased confidentiality and opportunistic encryption (e.g. [57]). Individuals have objected to the “anointing” of certain security technologies by these expert groups.¹⁵ This kind of pushback can be a reminder for the basic humility of voluntary standard-setting: consensus standards only have impact when widely implemented.

VI. FUTURE WORK

Improving privacy reviews for future Internet and Web standards is a matter of ongoing work by many individuals and organizations. But understanding how values like privacy and security are enacted in foundational Internet standards is of particular importance now, as standard-setting organizations and the larger engineering community respond to intelligence agencies' efforts to subvert security standards.

This paper has presented initial results from an ongoing research project. The broader goal is to consider the practices that affect privacy within technical standard-setting organizations for the Internet and the Web through a multi-modal, ethnographic approach: qualitative methods including further semi-structured interviews of the diversity of participants and stakeholders; automated and manual text analysis; and field notes from observing and participating in these processes.

The context of multistakeholder technical standard-setting provides a view of the challenges implementing privacy-by-design without formal methods of hierarchical control. What we learn about implementing privacy within a standard-setting process may also apply to areas with similarly collaborative characteristics or non-hierarchical organizational structures: open source software development, Internet governance bodies and other multistakeholder institutions.

ACKNOWLEDGMENT

Thanks to research participants for their informative and engaging interviews; the Telecommunications Policy Research Conference attendees for useful feedback on an earlier poster; Professors Deirdre Mulligan and Paul Duguid for their reviews;

¹⁵Discussions on the TAG mailing list at the end of last year show various concerns: lists.w3.org/Archives/Public/www-tag/2014Dec

IWPE reviewers for their detailed readings and comments; and the School of Information's Doctoral Research and Theory Workshop for productive, candid conversation.

Research used in this paper was supported in part by the U.S. Department of Homeland Security under Grant Award Number 2006-CS-001-000001, under the auspices of the Institute for Information Infrastructure Protection (I3P) research program.

REFERENCES

- [1] N. Doty and D. K. Mulligan, "Internet Multistakeholder Processes and Techno-Policy Standards: Initial Reflections on Privacy at the World Wide Web Consortium," *Journal on Telecommunications and High Technology Law*, vol. 11, 2013 [Online]. Available: http://www.jthtl.org/content/articles/V11i1/JTHTLv11i1_MulliganDoty.PDF
- [2] W. B. Gallie, "Essentially Contested Concepts," *Proceedings of the Aristotelian Society*, vol. 56, pp. 167–198, Jan. 1956 [Online]. Available: <http://www.jstor.org/stable/4544562>
- [3] D. K. Mulligan and C. Koopman, "Theorizing Privacy's Contestability: A Multi-Dimensional Analytic of Privacy," 2015.
- [4] WHATWG, "FAQ - WHATWG Wiki." Feb-2015 [Online]. Available: <https://wiki.whatwg.org/wiki/FAQ>. [Accessed: 04-Mar-2015]
- [5] P. Ford, "On HTML5 and the Group That Rules the Web," *The New Yorker*, Nov. 2014 [Online]. Available: <http://www.newyorker.com/tech/elements/group-rules-web>
- [6] OASIS, "About Us | OASIS." [Online]. Available: <https://www.oasis-open.org/org>. [Accessed: 04-Mar-2015]
- [7] J. Sabo, M. Willett, P. Brown, and D. Jutla, "Privacy Management Reference Model and Methodology (PMRM) V1.0," OASIS, technical report, Mar. 2012 [Online]. Available: <http://docs.oasis-open.org/pmr/PMRM/v1.0/csd01/PMRM-v1.0-csd01.html>
- [8] OASIS, "OASIS Privacy by Design Documentation for Software Engineers (PbD-SE) Technical Committee." [Online]. Available: https://www.oasis-open.org/committees/tc_home.php?wg_abbrev=pbds. [Accessed: 04-Mar-2015]
- [9] "About | Kantara Initiative." [Online]. Available: <https://kantarainitiative.org/about/>. [Accessed: 04-Mar-2015]
- [10] Ecma International, "Welcome to Ecma International." [Online]. Available: <http://www.ecma-international.org/>. [Accessed: 04-Mar-2015]
- [11] International Telecommunications Union, "About." [Online]. Available: <http://www.itu.int/en/about/Pages/default.aspx>. [Accessed: 04-Mar-2015]
- [12] S. Braman, "Privacy by design: Networked computing, 1969–1979," *New Media & Society*, vol. 14, no. 5, pp. 798–814, Aug. 2012 [Online]. Available: <http://nms.sagepub.com/content/14/5/798>
- [13] S. S. Shapiro, "Privacy by Design: Moving from Art to Practice," *Commun. ACM*, vol. 53, no. 6, pp. 27–29, Jun. 2010 [Online]. Available: <http://doi.acm.org/10.1145/1743546.1743559>
- [14] H. Flanagan and S. Ginoza, "RFC Style Guide," RFC Editor; Internet Requests for Comments; RFC Editor, RFC 7322, Sep. 2014 [Online]. Available: <http://www.rfc-editor.org/rfc/rfc7322.txt>
- [15] E. Rescorla and B. Korver, "Guidelines for Writing RFC Text on Security Considerations," RFC Editor; Internet Requests for Comments; RFC Editor, RFC 3552, Jul. 2003 [Online]. Available: <http://tools.ietf.org/html/rfc3552>
- [16] A. Rabkin, N. Doty, and D. K. Mulligan, "Facilitate, don't mandate," 2010 [Online]. Available: http://www.iab.org/wp-content/IAB-uploads/2011/03/nick_doty.pdf
- [17] H. Levkowetz, "Idnits Tool." Nov-2014 [Online]. Available: <http://tools.ietf.org/tools/idnits/>
- [18] "Security Directorate Review Process." [Online]. Available: <http://trac.tools.ietf.org/area/sec/trac/wiki/SecDirReview>. [Accessed: 01-Mar-2015]
- [19] A. Cooper, H. Tschofenig, B. Adoba, J. Peterson, J. Morris, M. Hansen, and R. Smith, "Privacy Considerations for Internet Protocols," Internet Architecture Board; Internet Requests for Comments, RFC 6973, 2013 [Online]. Available: <http://tools.ietf.org/html/rfc6973>
- [20] Internet Engineering Task Force, "Geographic Location/Privacy (geopriv)." Nov-2014 [Online]. Available: <http://www.ietf.org/wg/concluded/geopriv.html>. [Accessed: 04-Mar-2015]
- [21] N. Doty, D. K. Mulligan, and E. Wilde, "Privacy Issues of the W3C Geolocation API," UC Berkeley, School of Information, technical report, Feb. 2010 [Online]. Available: <http://escholarship.org/uc/item/0rp834wf>
- [22] S. Barocas, "Parsing Privacy Policies." [Online]. Available: http://solon.barocas.org/?page_id=200. [Accessed: 01-Mar-2015]
- [23] National Telecommunications & Information Administration, "Privacy Multistakeholder Process: Mobile Application Transparency." Nov-2013 [Online]. Available: <http://www.ntia.doc.gov/other-publication/2013/privacy-multistakeholder-process-mobile-application-transparency>. [Accessed: 01-Mar-2015]
- [24] L. Cranor, J. Reagle, J. Mackie-Mason, and D. Waterman, "Designing a Social Protocol: Lessons Learned from the Platform for Privacy Preferences Project," Lawrence Erlbaum Associates, 1998.
- [25] D. D. Clark, J. Wroclawski, K. R. Sollins, and R. Braden, "Tussle in cyberspace: defining tomorrow's Internet," *IEEE/ACM Transactions on Networking*, vol. 13, no. 3, pp. 462–475, Jun. 2005 [Online]. Available: http://ieeexplore.ieee.org/xpls/abs_all.jsp?arnumber=1458757
- [26] A. Schwartz, "Looking Back at P3P: Lessons for the Future," technical report November, 2009 [Online]. Available: https://www.cdt.org/files/pdfs/P3P_Retro_Final_0.pdf
- [27] FTC, "Protecting Consumer Privacy in an Era of Rapid Change: A Proposed Framework for Businesses and Policymakers," technical report December, 2010 [Online]. Available: <http://www.ftc.gov/reports/preliminary-ftc-staff-report-protecting-consumer-privacy-era-rapid-change-proposed-framework>
- [28] Electronic Frontier Foundation, "A privacy-friendly DNT Policy." 2014 [Online]. Available: <https://www.eff.org/dnt-policy>. [Accessed: 01-Mar-2015]
- [29] L. Mastria, "DAA Leaves W3C Tracking Protection Working Group To Convene A New Process on Browser-Based Signals and Consumer Privacy." Sep-2013 [Online]. Available: <http://www.aboutads.info/blog/daa-leaves-w3c->

- tracking-protection-working-group-convene-new-process-browser-based-signals-and-. [Accessed: 15-Jan-2015]
- [30] N. Doty, H. West, J. Brookman, S. Harvey, and E. Newland, "Tracking Compliance and Scope," W3C, Working Draft, Nov. 2014 [Online]. Available: <http://www.w3.org/TR/2014/WD-tracking-compliance-20141125/>
- [31] G. Greenwald, "NSA collecting phone records of millions of Verizon customers daily," *The Guardian*. Jun-2013 [Online]. Available: <http://www.theguardian.com/world/2013/jun/06/nsa-phone-records-verizon-court-order>
- [32] G. Greenwald and E. MacAskill, "NSA Prism program taps in to user data of Apple, Google and others," *The Guardian*. Jun-2013 [Online]. Available: <http://www.theguardian.com/world/2013/jun/06/us-tech-giants-nsa-data>
- [33] G. Greenwald, "XKeyscore: NSA tool collects 'nearly everything a user does on the internet'," *The Guardian*. Jul-2013 [Online]. Available: <http://www.theguardian.com/world/2013/jul/31/nsa-top-secret-program-online-data>
- [34] N. P. J. Larson and S. Shane, "N.S.A. Able to Foil Basic Safeguards of Privacy on Web," *The New York Times*. Sep-2013 [Online]. Available: <http://www.nytimes.com/2013/09/06/us/nsa-foils-much-internet-encryption.html>
- [35] M. Thomson, "A Simple Statement," IETF Secretariat, Internet-Draft draft-thomson-perpass-statement-00, Nov. 2013 [Online]. Available: <http://www.ietf.org/internet-drafts/draft-thomson-perpass-statement-00.txt>
- [36] Internet Engineering Task Force, "IETF 88 Registration System Attendance List." 2013 [Online]. Available: <https://www.ietf.org/registration/ietf88/attendance.py>. [Accessed: 01-Mar-2015]
- [37] R. Housley, "IETF88 Technical Plenary hums." Nov-2013 [Online]. Available: <http://www.ietf.org/mail-archive/web/ietf/current/msg83857.html>
- [38] S. Farrell and H. Tschofenig, "Pervasive Monitoring Is an Attack," RFC Editor; Internet Requests for Comments; RFC Editor, RFC 7258, May 2014 [Online]. Available: <http://tools.ietf.org/html/rfc7258>
- [39] "Perpass 'BoF' session - Considering pervasive monitoring," *IETF meeting agendas*. Nov-2013 [Online]. Available: <https://datatracker.ietf.org/meeting/88/agenda/perpass>. [Accessed: 01-Mar-2015]
- [40] S. Farrell, R. Wenning, B. Bos, M. Blanchet, and H. Tschofenig, "STRINT workshop report," IETF Secretariat; Working Draft, Internet-Draft draft-iab-strint-report-00, Apr. 2014 [Online]. Available: <http://www.ietf.org/internet-drafts/draft-iab-strint-report-00.txt>
- [41] C. Timberg, "Google encrypts data amid backlash against NSA spying," *The Washington Post*. Sep-2013 [Online]. Available: http://www.washingtonpost.com/business/technology/google-encrypts-data-amid-backlash-against-nsa-spying/2013/09/06/9acc3c20-1722-11e3-a2ec-b47e45e6f8ef_story.html
- [42] B. Gellman and A. Soltani, "NSA infiltrates links to Yahoo, Google data centers worldwide, Snowden documents say," *The Washington Post*. Oct-2013 [Online]. Available: http://www.washingtonpost.com/world/national-security/nsa-infiltrates-links-to-yahoo-google-data-centers-worldwide-snowden-documents-say/2013/10/30/e51d661e-4166-11e3-8b74-d89d714ca4dd_story.html
- [43] M. Nottingham, "Securing the Web," W3C Technical Architecture Group, W3C TAG Finding, Jan. 2015 [Online]. Available: <https://w3ctag.github.io/web-https/>
- [44] Internet Architecture Board, "IAB Statement on Internet Confidentiality." 2014 [Online]. Available: <https://www.iab.org/2014/11/14/iab-statement-on-internet-confidentiality/>. [Accessed: 02-Feb-2015]
- [45] IETF HTTP Working Group, "HTTP/2 Frequently Asked Questions." [Online]. Available: <http://http2.github.io/faq/#does-http2-require-encryption>. [Accessed: 01-Mar-2015]
- [46] M. West and Y. Zhu, "Privileged Contexts," W3C, Editor's Draft, Feb. 2015 [Online]. Available: <http://w3c.github.io/webappsec/specs/powerfulfeatures/>
- [47] H. Tschofenig, "IAB Privacy Consideration Tutorial Slides." Jul-2013 [Online]. Available: https://www.iab.org/?attachment_id=7038
- [48] M. West, "Strawman Self-Review Questionnaire: Security and Privacy," W3C, Unofficial Draft, Jan. 2015 [Online]. Available: <https://mikewest.github.io/spec-questionnaire/security-privacy/>
- [49] N. Doty, "Fingerprinting Guidance for Web Specification Authors," W3C, Unofficial Draft, Oct. 2014 [Online]. Available: <http://w3c.github.io/fingerprinting-guidance/>
- [50] F. Dawson, "Specification Privacy Assessment (SPA): Creating Privacy Considerations for W3C Technical Specifications," W3C, Unofficial Draft, Jun. 2013 [Online]. Available: <https://yrlesru.github.io/SPA/>
- [51] M. C. Oetzel and S. Spiekermann, "A systematic methodology for privacy impact assessments: a design science approach," *European Journal of Information Systems*, Jul. 2013 [Online]. Available: <http://www.palgrave-journals.com/ejis/journal/vaop/ncurrent/abs/ejis201318a.html>
- [52] S. Turner, "[privacydir] Closing ML." May-2012 [Online]. Available: <http://www.ietf.org/mail-archive/web/privacydir/current/msg00066.html>
- [53] Privacy Interest Group, "Privacy Interest Group Teleconference - 04 Dec 2014." Dec-2014 [Online]. Available: <http://www.w3.org/2014/12/04-privacy-minutes.html>. [Accessed: 01-Mar-2015]
- [54] Internet Architecture Board Privacy Program, "Privacy Reviews." [Online]. Available: <https://www.iab.org/activities/programs/concluded-programs/privacy-program/privacy-reviews/>. [Accessed: 01-Mar-2015]
- [55] S. Farrell, "[perpass] privacy/PM reviews of existing stuff." Jan-2014 [Online]. Available: <http://www.ietf.org/mail-archive/web/perpass/current/msg01530.html>
- [56] P. Hoffman, "The Tao of IETF: A Novice's Guide to the Internet Engineering Task Force." Aug-2012 [Online]. Available: <https://www.ietf.org/tao.html>. [Accessed: 03-Feb-2015]
- [57] M. Nottingham and M. Thomson, "Opportunistic Security for HTTP," IETF Secretariat; Working Draft, Internet-Draft draft-ietf-httpbis-http2-encryption-01, Dec. 2014 [Online]. Available: <http://www.ietf.org/internet-drafts/draft-ietf-httpbis-http2-encryption-01.txt>