



Response to Department of Commerce's «Commercial Data Privacy and Innovation in the Internet Economy: A Dynamic Policy Framework»

The Department of Commerce's Internet Policy Task Force is examining policy approaches that reduce barriers to digital commerce while strengthening protections for commercial data privacy, cybersecurity, intellectual property, and the global free flow of information. Within the published Green Paper, the IPTF asks for specific feedback on various aspects of the overall depicted framework.

- I. Introduction
- II. Suggestions concerning DoC Questions

I. Introduction

W3C is an international community where Member organizations, a full-time staff, and the public work together to develop Web standards. Led by Web inventor Tim Berners-Lee and CEO Jeffrey Jaffe, W3C's mission is to lead the Web to its full potential.

W3C efforts related to privacy on the Web began in 1997, when development started on the widely known Platform for Privacy Preferences (P3P), published as a Web Standard in 2002.

The W3C staff have been part of the broader privacy conversation throughout the last decade, and have participated in many different research projects on Privacy in the United States and in Europe, including the Transparent Accountable Datamining Initiative, Policy Aware Web, Theory and Practice of Accountable Systems, PRIME and PrimeLife. One important vehicle for making connections from research work to other work is the W3C Policy Languages Interest Group (PLING), which also helps to bridge communities fragmented around policy languages and access control. Findings from the research influenced the work carried out in other W3C Working Groups, but not to the extent we had hoped for.

The role of the standards W3C builds is increasingly broad: W3C is no longer tied to the document mindset of the early Web; instead, we build the standardized underpinnings for what looks increasingly like a Web operating system: General purpose data formats, general purpose communications frameworks, general purpose APIs that make device features accessible to the Web that had previously been outside the sandbox.

As we build and design advanced APIs that permit access to risky features, topics like the transparency of the data collection itself, limiting the scope of user errors, or the user's ability to recover from erroneously granted consent take center stage. These factors at times influence the design of APIs (does the user pass a selection of cards from his address book to a web site, or does he grant the web site access to the address book). At other times, all we might be able to do in specifications is to sketch basic requirements, as the distinction between a privacy friendly and a dangerous implementation may be entirely dependent on the details of user interfaces and interactions, beyond the scope of what can be reasonably specified.

II. Suggestions concerning DoC Questions

1. The Task Force recommends adoption of a baseline commercial data privacy framework built on an expanded set of Fair Information Practice Principles (FIPPs).
 - a. Should baseline commercial data privacy principles, such as comprehensive FIPPs, be enacted by statute or through other formal means to address how current privacy law is enforced?

Answer

Establishing baseline commercial data privacy principles would contribute to the further harmonization of the global e-commerce market at least for the countries attached to the OECD, and improve the transatlantic relations on online services of all sorts. It would also allow to remedy the most blatant and harmful abuses of privacy across all types of services and thus contribute to more trust and peace in the marketplace, and increase consumer confidence.

-
- b. How should baseline privacy principles be enforced? Should they be enforced by non-governmental entities in addition to being the basis for FTC enforcement actions?
 - c. As policymakers consider baseline commercial data privacy legislation, should they seek to grant the FTC the authority to issue more detailed rules? What criteria are useful for deciding which FIPPs require further specification through rule making under the Administrative Procedure Act?

d. Should baseline commercial data privacy legislation include a private right of action?

Answer

The answers to b., c. and d. are related: The European system treats enforcement as a public function and gives citizens the possibility to complain to data commissioners. The US equivalent would be a system that puts the responsibility for enforcement into the hands of the FTC. As noted in W3C's response to the Notice of Inquiry, such a system won't really scale to the volume of personal information bearing interaction on the Web. However, it would contribute to addressing the most blatant abuses.

Individual citizens, despite having rights, can't actually enforce those rights when it is very difficult to establish an actual monetary damage for most privacy violations -- leaving severe identity theft cases aside. Researchers believe that only a system that gives citizens both rights and standing to sue, and introduces some kind of minimum damage to privacy breaches will allow the legal system to apply its usual scaling mechanisms (such as class action suits) to the privacy field. Carefully crafted rights for citizens will thus enhance the legal enforcement of privacy rights. The obvious risk of this approach is that it might introduce an unhealthy amount of litigation into the online marketplace, in particular given the complexity inherent in framing the privacy aspects of even every-day data processing online.

The right to access and rectify one's own data needs further attention. There is a twofold technical challenge in bringing subject access into the Internet age: The first challenge is that any mechanism that permits the data subject to access one's own personal data has to make ensure that that the data access is limited to that subject. In addition to this authentication and access control challenge, common protocols for subject access need to be defined that permit software to communicate and inter-operate and addresses, and enable comprehensible and simple interactions for the user.

-
2. To meet the unique challenges of information intensive environments, FIPPs regarding enhancing transparency; encouraging greater detail in purpose specifications and use limitations; and fostering the development of verifiable evaluation and accountability should receive high priority.
- a. What is the best way of promoting transparency so as to promote informed choices? The Task Force is especially interested in comments that address the benefits and drawbacks of legislative, regulatory, and voluntary private sector approaches to promoting transparency.

Answer

As we noted in our response to the Notice of Inquiry, a truly informed choice requires user understanding. The complexity of our society and of the commercial relations is such that even if users have all the relevant information available, they typically won't be able to grasp what this information actually means for them. Communicating privacy decisions and

circumstances in a meaningful and usable way is a relatively recent area of active research and development. Further research is needed, not only on user-machine interfaces that take users' cognitive abilities and limitations into account, but also on users' social and philosophical expectations. Web science suggests that there is a difficult interdisciplinary field to explore where social, technical, cognitive, and psychological issues mix.

- b. What incentives could be provided to encourage the development and adoption of practical mechanisms to protect consumer privacy, such as PIAs, to bring about clearer descriptions of an organization's data collection, use, and disclosure practices?
 - c. What are the elements of a meaningful PIA in the commercial context? Who should define these elements?
 - d. What processes and information would be useful to assess whether PIAs are effective in helping companies to identify, evaluate, and address commercial data privacy issues?
 - e. Should there be a requirement to publish PIAs in a standardized and/or machine-readable format?
-

Answer

While b. - e. are concrete questions that almost suggest a Yes/No answer, we think that *"Privacy Impact Assessments (PIA)"* raise a variety of challenges and opportunities that are interlinked with the entire system.

In 2002, many critics of P3P pointed out the allegedly low impact of labeling one's offer on the Web. The experience and feedback of W3C shows that expressing a privacy policy required businesses to find out what they were actually doing with all the data they collected, and whether the data was really needed. PIAs of any form will only be successful if there is a tangible advantage that outweighs the burden to create them.

Other factors also affect the consideration of PIAs. In the framework of P3P, W3C worked 2001-2003 together with the Comité Européen de Normalisation (CEN) on a Workshop Agreement called IPSE. This was supposed to produce standard contract clauses that could be attached to contracts as requirements and that could be audited. The IT industry fiercely opposed this approach, especially fighting against any form of external audit. The external audit was described as likely to *"waste high amounts of money without any tangible outcome"*. From his experience, it appears likely that the IT industry will favor a system of PIAs based on self assessment.

Today's privacy policies already imply some kind of self assessment. They are typically developed by legal departments eager to minimize companies' legal liabilities. This goal is best met by statements that produce maximum legal uncertainty (to not shy customers away), but still avoid concrete promises that companies could be held liable for. This is a consequence of the incentives set by an enforcement system that focuses on keeping online actors to holding their, promises but does not mandate any baseline practices.

In corporate decision-making, a possibly vague promise of better marketing

and user trust can be weak compared to the general counsel's prediction of significant liability. A concrete tangible potential harm is traded against some unproven promise. Therefore, how a PIA system alters the liability and enforcement environment for privacy breaches will be a key factor in its success. If the system grows the liability, incentives are aligned against companies putting themselves on the forefront of those setting up PIAs.

A machine readable format can play several roles for PIAs. Machine-readable formats don't carry the same level of nuance and detail as legal language: They tend to force precision into the system. Machine-readable formats also permit for fast search and matching of specific things that a consumer may not find acceptable. Thus proscribing a machine readable format can be an effective cure against large and imprecise declarations. However, the precision of machine-readable statements and their possible divergence from underlying legal language can also lead to divergence, and in turn increase the risk of liability. These caveats have to be either accepted or taken into account if proscribing a machine readable PIA is intended.

A machine readable format has another advantage. The Internet industry is collecting large amounts of personal data. Once collected, the data can later be mined and processed based on all kinds of aspects, and for all sorts of purposes. Promises made to the user, limitations implied by the context of collection (e.g., religious convictions) are not transported and data can be re-used to the detriment of the data subject, e.g. by insurances that might increase rates based on socially common data sharing on social networks. The issue facing the user is that data from one context is reappearing in another. If data collection is associated with machine readable metadata on context and conditions of the collection, context and conditions can be preserved in a machine readable way, and transported as the data is flowing through the system. The context and promises made become a *"sticky policy"* also able to bind all commercial use of that data downstream. It allows for privacy properties being taken into account by data mining algorithms. It facilitates machine mediated auditing of processing- and access decisions.

So exposing a machine readable PIA to the public at collection time has two benefits: It informs the user and it allows businesses to treat the data right in the subsequent processing. This advantage has to be weighed against a possible increased liability risk, and the resulting disincentives.

-
- f. What are consumers' and companies' experiences with systems that display information about companies' privacy practices in contexts other than privacy policies?
 - g. What are the relative advantages and disadvantages of different transparency-enhancing techniques in an online world that typically involves multiple sources being presented through a single user interface?

Answer:

Transparency mechanisms need to provide users with comprehensible, clear and actionable information that can be perceived, but does not unnecessarily

distract users from the task at hand.

User agents can contribute to transparency mechanisms where they can present information that fits into usage and interaction patterns. We encourage this approach where possible, but recognize that relatively simple and coherent models are needed in order to enable it.

In more complex cases, the responsibility for transparency about data collection, processing and usage cannot be placed on user agents alone: Instead, it is web applications' responsibility to ensure transparency in a way that integrates useably with users' decision and interaction flow. We acknowledge that this responsibility can become complex in multi-source scenarios. In cases where multiple applications and sources are routinely mixed, the development of common practices, possibly augmented with machine-readable policy expression, might help to make transparency more feasible.

We further observe that classical privacy policies, expressed in legal language, have proven ineffective in achieving actual transparency. An approach that requires users to pursue links to every single participant in a multi-source application in order to learn about this participant's privacy practices strikes us as similarly ineffective.

-
- h. Do these (dis)advantages change when one considers the increasing use of devices with more limited user interface options?

Answer

Mobile device constraints and usage modes reinforce the considerations that we have given in our previous response.

-
- i. Are purpose specifications a necessary or important method for protecting commercial privacy?

Answer

Expressing a purpose is expressing a use limitation in a whitelist approach. The EU approach to data protection is based on a prohibition of the processing of personal data with exceptions. Preventing the reuse of the data left at the doctor's to calculate the insurance rates is a good example where purpose and use limitations come together and are virtually the same.

From a technical perspective, the virtually indefinite number of possible purposes poses obvious issues. For commerce, one might come up with a definition of 30-50 purposes that would cover 80% of commercial transactions -- and would thus limit the use of acquired data to fulfill that purpose. This model generates question on how exact a purpose specification can be, or

whether it encompasses a variety of actions. The semantics of the defined boundaries are not carved in stone yet.

From a legal perspective, the advantage of a purpose-based system is that it can specify a given set of possible purposes without opening data up for undetermined use. In contrast, a use limitation essentially specifies a blacklist of prohibited purposes, and therefore allows all others.

So while purpose limitations may make sense in a legal setup, they are very difficult to follow and express in a technical setup.

-
- j. Currently, how common are purpose specification clauses in commercial privacy policies?
 - k. Do industry best practices concerning purpose specification and use limitations exist? If not, how could their development be encouraged?

Answer

The P3P 1.1 W3C Working Group Note contains a set of commonly used commercial purposes. The list could be enhanced based on additional quantitative research, based on data and privacy policies available on the Web today.

-
- l. What incentives could be provided to encourage companies to state clear, specific purposes for using personal information?

Answer

Finding incentives for companies in this area is a tough task, as it is inevitably tied to the liability issue. As we noted before, a system based on the enforcement of promises that companies can freely give will encourage companies to make the weakest promises possible, and express the broadest possible purpose to achieve the tiniest use limitations.

The second issue this will be tied to is purpose changes below.

-
- m. How should purpose specifications be implemented and enforced?

Answer

We take this question from a technical point of view as questions of legal enforcement were already discussed above. As we noted, purpose specifications are whitelist-like use limitations in a technical system. This means that software knowing the semantics for a limited set of primary and secondary purposes can monitor access to and flow of information according to the purpose rules set with that data. This technology is not mature yet, but

proof of concept and scientific implementations have been made and tested. Dealing with privacy as a metadata problem of the backend is a promising path to better reliability and trust for consumers and companies (liability) alike.

- n. How can purpose specifications and use limitations be changed to meet changing circumstances?

Answer

We believe that mechanisms that can renegotiate purpose and use limitations with the user in a privacy-friendly and non-surprising way are a promising area of research and development.

- o. Who should be responsible for demonstrating that a private sector organization's data use is consistent with its obligations? What steps should be taken if inconsistencies are found?

Answer

At the W3C Workshop on Privacy and data usage control, Jacques Bus, who headed the EU's security research funding administration for years, was calling for liability rules for Privacy and Security breaches. But the function of liability is more complex than just the sanctioning of unwanted behavior. There is first the challenge to find inconsistencies and there are many reasons for inconsistencies that have nothing to do with bad faith. In systems that follow the system of machine readable privacy metadata it is much easier to ensure consistency, but also much easier to have instant computer mediated audits to discover inconsistencies.

- p. Are technologies available to allow consumers to verify that their personal information is used in ways that are consistent with their expectations?

Answer

Most of the scandals that surfaced in the Privacy arena so far have been discovered without special verification rights. Europe has subject access. This is good for citizens and consumers to know what is held about them and to develop some trust. But subject access was rarely the means to uncover abusive privacy practices. It is also rare that user interactions for subject access are easily available to end users, and comprehensible.

As a necessary precondition for strong user-facing accountability, companies would need to deploy technologies that permit an audit of the circumstances and context under which data were collected, promises given, and any

processing of data that has occurred. These technologies will also require non-trivial interoperability properties.

- q. Are technologies available to help companies monitor their data use, to support internal accountability mechanisms?
-

Answer

As mentioned above, it is not sufficient to have database software and technologies to attach metadata. One of the bigger issues is the interoperability around the semantics of the privacy metadata. A frequent effect in recent research projects has been that projects started exploring new terms, but ended up reusing the P3P statement vocabulary. It would be interesting to hear from the industry, which semantics are missing.

- r. How should performance against stated policies and practices be assessed?
-

Answer

Performance is hard to measure in Privacy. Is the question to pinpoint a state on the scale from anonymity to full transparency of the user? Ultimately, performance assessment can't be constantly monitored and will be discovered in the usual political process. In really critical areas of information processing, e.g. concerning especially sensitive data, like religion, political opinions etc, there may be a closer monitoring needed.

- s. What incentives could be provided to encourage companies to adopt technologies that would facilitate audits of information use against the company's stated purposes and use limitations?
-

Answer

What if people who care (but make mistakes) will have lower liability than people who do not care? The incentive to install a complex system allowing a data governance by privacy metadata is to avoid incidents and to not trouble the normal course of business. Such a system, because it is based on metadata, has much more semantics and is much easier and cheaper to audit. So a combination of liability advantages and ease of governance may be sufficient to justify the pain of upgrading the systems.

3. Voluntary, enforceable codes of conduct should address emerging technologies and issues not covered by current application of baseline FIPPs. To encourage the development of such codes, the Administration should consider a variety of options, including (a) public statements of Administration support; (b) stepped up

FTC enforcement; and (c) legislation that would create a safe harbor for companies that adhere to appropriate voluntary, enforceable codes of conduct that have been developed through open, multi-stakeholder processes.

Answer

W3C welcomes this statement as it makes investments predictable. Part of the issue with Privacy Enhancing Technologies (PET) in the past 10 years was that regulators were waiting for technology to be developed and would judge technology afterwards. Development of technology has a high cost. At the same time, privacy wasn't seen as a cause of dramatic increase in revenues. Privacy was seen as a good housekeeping exercise for businesses. In this context, the necessary investment is only justifiable if the authorities hint at some benefit at the end of the development process. As the authorities refused for the sake of independence, the investment of industry was rather low. PET research was more a matter of public research. With a commitment from the regulators, this landscape can be changed.

-
4. Using existing resources, the Commerce Department should establish a Privacy Policy Office (PPO) to serve as a center of commercial data privacy expertise. The proposed PPO would have the authority to convene multi-stakeholder discussions of commercial data privacy implementation models, best practices, codes of conduct, and other areas that would benefit from bringing stakeholders together; and it would work in concert with the Executive Office of the President as the Administration's lead on international outreach on commercial data privacy policy. The PPO would be a peer of other Administration offices and components that have data privacy responsibilities; but, because the PPO would focus solely on commercial data privacy, its functions would not overlap with existing Administration offices. Nor would the PPO would have any enforcement authority.
- a. Should the FTC be given rulemaking authority triggered by failure of a multi-stakeholder process to produce a voluntary enforceable code within a specified time period?

Answer

Tracking technology uses the Web. This technology is global. Codes of conduct and other rulemaking on the edge between technology and regulation has global effects. Thus we welcome the intention to work out codes of conduct in a multi-stakeholder process. As this process affects the technical layers of the global Internet and Web, the participation of technical stakeholders such as global Internet standardization organizations will be a critical element of the global stakeholder concertation.

-
- b. How can the Commerce Department best encourage the discussion and development of technologies such as «Do Not Track»?

Answer

The Commerce Department's leadership in the space have already triggered new enthusiasm for the Web privacy discussion. We encourage the Commerce Department to continue on this way, and recommend connecting the ongoing discussion with technical, regulatory and legal conversations worldwide.

- c. Under what circumstances should the PPO recommend to the Administration that new policies are needed to address failure by a multi-stakeholder process to produce an approved code of conduct?
 - d. How can cooperation be fostered between the National Association of Attorneys General, or similar entities, and the PPO?
5. The FTC should remain the lead consumer privacy enforcement agency for the U.S. Government.
 - a. Do FIPPs require further regulatory elaboration to enforce, or are they sufficient on their own?
 - b. What should be the scope of FTC rulemaking authority?
 - c. Should FIPPs be considered an independent basis for FTC enforcement, or should FTC privacy investigations still be conducted under Federal Trade Commission Act Section 5 «unfair and deceptive» jurisdiction, buttressed by the explicit articulation of the FIPPs?
 - d. Should non-governmental entities supplement FTC enforcement of voluntary codes?
 - e. At what point in the development of a voluntary, enforceable code of conduct should the FTC review it for approval? Potential options include providing an ex ante «seal of approval,» delaying approval until the code is in use for a specific amount of time, and delaying approval until enforcement action is taken against the code.
 - f. What steps or conditions are necessary to make a company's commitment to follow a code of conduct enforceable?
6. The U.S. government should continue to work toward increased cooperation among privacy enforcement authorities around the world and develop a framework for mutual recognition of other countries' commercial data privacy frameworks. The United States should also continue to support the APEC Data Privacy Pathfinder project as a model for the kinds of principles that could be adopted by groups of countries with common values but sometimes diverging privacy legal frameworks.
7. Consideration should be given to a comprehensive commercial data security breach framework for electronic records that includes notification provisions, encourages companies to implement strict data security protocols, and allows States to build upon the framework in limited ways. Such a framework should track the effective protections that have emerged from State security breach notification laws and policies.
 - a. What factors should breach notification be predicated upon (e.g., a risk assessment of the potential harm from the breach, a specific threshold such as number of records, etc.)?
8. A baseline commercial data privacy framework should not conflict with the strong sectoral laws and policies that already provide important protections to Americans, but rather should act in concert with these protections.
 - a. Are there lessons from sector-specific commercial data privacy laws—their development, their contents, or their enforcement—that could inform general U.S. commercial data privacy policy?
9. Any new Federal privacy framework should seek to balance the desire to create uniformity and predictability across State jurisdictions with the desire to permit States the freedom to protect consumers and to regulate new concerns that arise from emerging technologies, should those developments create the need for

additional protection under Federal law.

- a. Should a preemption provision of national FIPPs-based commercial data privacy policy be narrowly tailored to apply to specific practices or subject matters, leaving States free to regulate new concerns that arise from emerging technologies? Or should national policy, in the case of legislation, contain a broad preemption provision?
 - b. How could a preemption provision ensure that Federal law is no less protective than existing State laws? What are useful criteria for comparatively assessing how protective different laws are?
 - c. To what extent should State Attorneys General be empowered to enforce national FIPPs-based commercial data privacy legislation?
 - d. Should national FIPPs-based commercial data privacy legislation preempt State unfair and deceptive trade practices laws?
10. The Administration should review the Electronic Communications Privacy Act (ECPA), with a view to addressing privacy protection in cloud computing and location-based services. A goal of this effort should be to ensure that, as technology and market conditions change, ECPA continues to appropriately protect individuals' expectations of privacy and effectively punish unlawful access to and disclosure of consumer data.
- a. The Task Force seeks case studies and statistics that provide evidence of concern—or comments explaining why concerns are unwarranted—about cloud computing data privacy and security in the commercial context. We also seek data that link any such concerns to decisions to adopt, or refrain from adopting, cloud computing services.
 - b. The Task Force also seeks input on whether the current legal protections for transactional information and location information raise questions about what privacy expectations are reasonable and whether additional protections should be mandated by law. The Task Force also invites comments that discuss whether privacy protections for access to location information need clarification in order to facilitate the development, deployment and widespread adoption of new location-based services.

Answer

After long and intense discussions, W3C is about to established a Geolocation API Specification for location based services on the web. It has to be noted that no agreement could be achieved towards a normative requirement for the visualization of an active geolocation sharing as it is forseen in Art. 9 of Directive 2002/58EC. Nevertheless, most of the widespread implementations of the Geolocation API Specification have an indicator that signals if the geolocation API is active, and allows for easy revocation. For further issues and important rules, the section 4 “*Security and privacy considerations*” is reproduced here:

4 Security and privacy considerations

The API defined in this specification is used to retrieve the geographic location of a hosting device. In almost all cases, this information also discloses the location of the user of the device, thereby potentially compromising the user's privacy. A conforming implementation of this specification must provide a mechanism that protects the user's privacy and this mechanism should ensure that

no location information is made available through this API without the user's express permission.

4.1 Privacy considerations for implementors of the Geolocation API

User agents must not send location information to Web sites without the express permission of the user. User agents must acquire permission through a user interface, unless they have prearranged trust relationships with users, as described below. The user interface must include the URI of the document origin [DOCUMENTORIGIN]. Those permissions that are acquired through the user interface and that are preserved beyond the current browsing session (i.e. beyond the time when the browsing context [BROWSINGCONTEXT] is navigated to another URL) must be revocable and user agents must respect revoked permissions.

Some user agents will have prearranged trust relationships that do not require such user interfaces. For example, while a Web browser will present a user interface when a Web site performs a geolocation request, a VOIP telephone may not present any user interface when using location information to perform an E911 function.

4.2 Privacy considerations for recipients of location information

Recipients must only request location information when necessary. Recipients must only use the location information for the task for which it was provided to them. Recipients must dispose of location information once that task is completed, unless expressly permitted to retain it by the user. Recipients must also take measures to protect this information against unauthorized access. If location information is stored, users should be allowed to update and delete this information.

The recipient of location information must not retransmit the location information without the user's express permission. Care should be taken when retransmitting and use of encryption is encouraged.

Recipients must clearly and conspicuously disclose the fact that they are collecting location data, the purpose for the collection, how long the data is retained, how the data is secured, how the data is shared if it is shared, how users may access, update and delete the data, and any other choices that users have with respect to the data. This disclosure must include an explanation of any exceptions to the guidelines listed above.

4.3 Additional implementation considerations

This section is non-normative.

Further to the requirements listed in the previous section, implementors of the Geolocation API are also advised to consider the following aspects that may negatively affect the privacy of their users: in certain cases, users may inadvertently grant permission to the user agent to disclose their location to Web sites. In other cases, the content hosted at a certain URL changes in such a way that the previously granted location permissions no longer apply as far as the user is concerned. Or the users might simply change their minds.

Predicting or preventing these situations is inherently difficult. Mitigation and in-depth defensive measures are an implementation responsibility and not prescribed by this specification. However, in designing these measures, implementors are advised to enable user awareness of location sharing, and to provide easy access to interfaces that enable revocation of permissions.

-
- c. The Task Force seeks information from the law enforcement community regarding the use of ECPA today and how investigations might be affected by proposed amendments to ECPA's provisions.

For questions about W3C or the answers here, please contact Rigo Wenning (rigo@w3.org)