

By Annie I. Antón,
Elisa Bertino, Ninghui Li,
and Ting Yu

A Roadmap for COMPREHENSIVE ONLINE PRIVACY POLICY MANAGEMENT

A framework supporting the privacy policy life cycle helps guide the kind of research to consider before sound privacy answers may be realized.

Information technology advances are making Internet and Web-based system use the common choice in many application domains, ranging from business to health care to scientific collaboration and distance learning. However, adoption is slowed by well-founded concerns about privacy, especially given that data collected about individuals is being combined with information from other sources and analyzed by powerful tools (such as data mining tools).

Effective solutions for privacy protection are of interest to industry, government, and society at large, but the challenge is to satisfy the often-conflicting requirements of all these stakeholders. Enterprises need mechanisms to ensure their systems are compliant with both the policies they articulate and the law. Moreover, they need to understand how to specify, deploy, communicate, and enforce privacy policies. Legislators and regulatory bodies require mechanisms to verify how privacy-related laws are actually enforced by enterprises in their software systems. Finally, end users must be able to easily understand privacy policies [2] and need effective, transparent, and comprehensible online privacy-protection mechanisms.

Significant efforts in industry are seeking to better protect sensitive information online and better communicate the mechanisms used to do so in the form of privacy policies. However, existing solutions are still fragmented and far from satisfactory. For example, existing languages for specifying privacy policies lack formal and unambiguous semantics, are limited in expressive power, and lack enforcement and auditing support [9]. End user privacy management tools are limited in capability or difficult to use. To provide effective online privacy protection, a comprehensive framework that covers the entire privacy policy life cycle is needed. This life cycle includes enterprise policy creation, enforcement, analysis, and auditing, as well as end user agent presentation and privacy policy processing. Trustworthy privacy protection can only be attained when broad consideration is given not only to IT solutions, but also to a wide range of perspectives from other disciplines. To this end, technical attempts to support privacy policy management must take into account the human, legal, and economic perspectives [1, 6] relevant to privacy. For example, how much are individuals willing to pay to keep their information private?

Here, we present a comprehensive architectural framework that supports the privacy policy life cycle. We identify the relevant technological and non-technical components required to support this life cycle, showing the relationships between these components. The framework suggests a detailed roadmap for research to be undertaken before sound privacy solutions may be realized.

PRIVACY POLICY TECHNOLOGIES

To make privacy policies more readable and enforceable, two privacy policy specification languages have emerged: P3P and EPAL.

Platform for Privacy Preferences (P3P) Project. The W3C's P3P Project [12] enables Web sites to

encode their data collection and data-use practices in a machine-readable XML format known as P3P policies. The W3C has also designed APPEL (a P3P Preference Exchange Language) that allows users to specify their privacy preferences. Ideally, through the use of P3P and APPEL, a user agent (a program working on the user's behalf) should be able to check a Web site's privacy policy against the user's privacy preferences, and automatically determine whether the Web site's data collection and data-usage practices are acceptable to the user. P3P appears to be the most widely used (if not the only) language for encoding enterprises' privacy policies for consumption by end users. However, P3P has several limitations and shortcomings that must be addressed.

The P3P language does not have a clear semantics and can therefore be interpreted and presented differently by different user agents. Companies may be reluctant to provide P3P policies on their Web sites because policies may be misrepresented. Quoting from CitiGroup's position paper [11], "The same P3P policy could be represented to users in ways that may be counter to each other as well as to the intent of the site...This results in legal and media risk for companies implementing P3P that needs to be addressed and resolved if P3P is to fulfill a very important need." Furthermore, a policy specified in P3P may be internally inconsistent [9].

The fundamental reason underlying the aforementioned technical difficulties is that the need for a semantics was apparently overlooked in the initial design of P3P, leaving too much freedom for user agents to misinterpret P3P policies. As discussed in [9], the problem is not just about the ambiguity of vocabulary in P3P, but also about how the different components (such as collected data items, purposes, recipients, and retentions) in a P3P statement interact. Additionally, the expressive power of P3P is limited [10, 11]. Many statements in a natural language privacy policy cannot be expressed in P3P, including, for example, how long data will be stored, what security mechanisms are in place to protect stored data, what kinds of data are not collected or shared, and so on.

Though Web sites are starting to post their P3P policies, the majority of online privacy policies are published in natural language. Currently, only textual policies are legally binding for an enterprise. Natural-language privacy policies cover a much broader scope of an enterprise's privacy practices than P3P policies. Moreover, natural-language policies tend to be more ambiguous and incomplete [2], making it difficult to maintain consistency between the policies and their more formal machine-readable representations. Through an iterative process, ambiguities can be

Privacy policies represent long-term promises made by an enterprise to its end users and are determined by business practice and legal concerns. It is undesirable to change an enterprise's promises to customers every time an internal access control rule changes.

removed; however, tools are needed for translating natural-language policies into machine readable and enforceable policies to facilitate consistency checking. Policy translation tools will enable large-scale processing of textual privacy policies and increase general understanding about the current state of privacy practices.

The P3P framework does not address enforcement or auditing. An enterprise has no way to determine whether published privacy policy is actually enforced within their information systems; nor can it prove to other parties that adequate procedures have been followed to ensure compliance with its privacy policy. This problem is exacerbated by the fact that an enterprise shares customer data with other business partners, which may have different privacy practices [3]. Even within a single organization, multiple privacy policies often exist [2]. Tools are thus needed for comparing and analyzing different privacy policies, and to enforce privacy-aware information flow to thwart inappropriate information flows [3].

Enterprise Privacy Policy Enforcement. Researchers at IBM are developing enterprise privacy architecture solutions. Karjoth et al. [8] proposed a privacy-centric access control language—E-P3P and its successor EPAL. EPAL (Enterprise Privacy Authorization Language) is an abstract-level access control language, with features devoted to privacy protection; for example, data accessing purposes. Additionally, IBM is developing a privacy logging and reporting framework that maps application-specific information to enterprise-level information [4]. We identify the following limitations of existing work.

First, the efficient and correct enforcement of policies specified in EPAL (or in a language for similar purposes) in the data storage layer has not been addressed. Policies specified at the EPAL level must be enforced at the time data is accessed. In most cases, such data is stored in databases and is accessed frequently. Thus, if every data access had to rely on external policy evaluation, the performance would be unacceptable.

Second, the relationship between policies at the P3P level and the EPAL level has not been adequately addressed. Karjoth et al. [7] proposed to generate P3P policies from EPAL policies. We disagree with this approach. Privacy policies represent long-term promises made by an enterprise to its end users and are determined by business practice and legal concerns. On the other hand, access control policies represent internal data handling practices that may change more frequently. It is undesirable to change an enterprise's promises to customers every time an internal access control rule changes. In fact, a privacy enforcement mechanism should be able to grandfather data and associated policies (to limit scope of impact when policies change).

Third, EPAL does not address situations arising from information flows between applications under different privacy policies. The sticky policy paradigm [8], which associates relevant consents with users' data so that they can be enforced during access control decisions, can help to a certain extent. However, most data exchange interfaces today do not support sticky policies; theory and tools to control information flows to other applications governed by different privacy policies are needed to ensure the correct privacy policy is enforced.

A COMPREHENSIVE FRAMEWORK FOR ONLINE PRIVACY PROTECTION

We now provide a general overview of the framework's key components and desirable functionalities and interactions. Figure 1 shows the architectural representation of a framework for privacy policy management.

Enterprise side. To support the complete life cycle of a privacy policy, the framework's enterprise side is organized according to a three-tier model.

Top tier (principles of privacy practices): An enterprise's high-level privacy promises are specified in privacy policies (using formal and/or natural language).

Although many organizations now post online privacy policies, these organizations must realize that simply posting a privacy policy on their Web site does not guarantee compliance with existing legislation.

Policies in this tier are intended for general Internet users. They should be specified by dedicated privacy officers who are familiar with both the enterprise's business practice and relevant privacy law and regulations. Key challenges include the design of a precise semantic model for privacy policies and expressive formal privacy policy languages. Policy languages for this tier should focus on which privacy goals are to be achieved, rather than how to achieve them.

Middle tier (security policies): In this layer, traditional security policies, for example, those governing authentication, access control, and information flow are needed to enforce high-level privacy policies. Policies at this tier should be specified by security officers who are familiar with high-level privacy policies and with the business processing needs of specific application domains. Within one application, privacy-centric access control and auditing policies ensure that data access does not violate privacy policies or security requirements. Data models and user management models are required to track how collected information is used by applications. A key challenge is the need to guarantee consistency between application-specific access control policies and privacy promises. Policy authoring and analysis tools based on specific application models are also required. Furthermore, because data may flow between applications governed by different high-level privacy policies, information flow control policies are needed to ensure that such data flow does not violate privacy promises.

The access control and auditing policies in this tier are application specific, but are usually independent of application implementation details. In an applica-

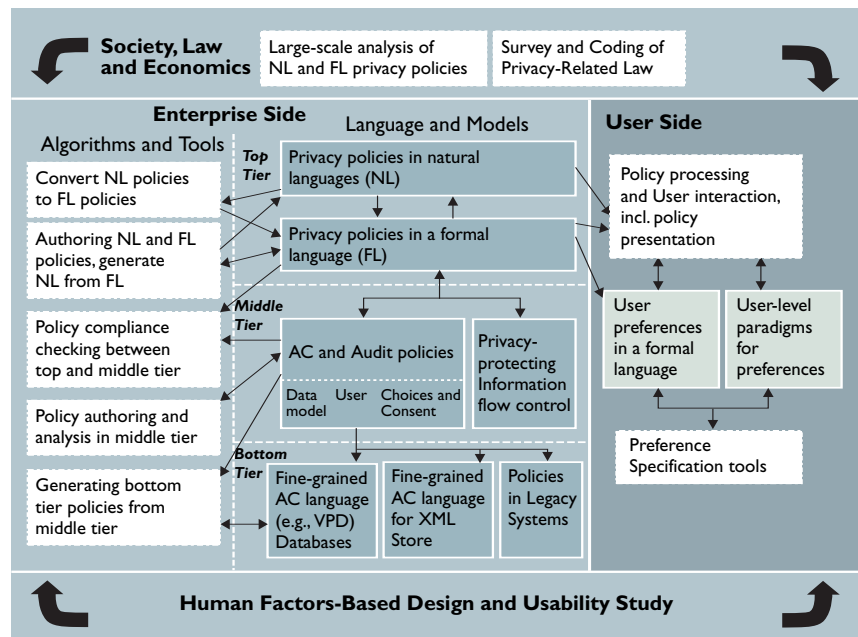


Figure 1. The architecture of a comprehensive framework for online privacy policy management.

tion, different levels of abstractions are commonly exploited to ease management overhead. For example, the model of information flow in an organization is usually independent of the physical storage of the information and the mechanisms through which information is exchanged between different departments. The separation of logical information flow and its physical storage and exchange implies the need for another level of privacy policy enforcement.

Bottom tier (enforcement in the physical layer): Access control and auditing policies must be materialize through policy configurations in the underlying information repository. The nature of privacy policies tends to be fine-grained, for example, each individual user may allow different usages of his or her data. Thus, fine-grained access control is needed; for exam-

ple, if relational databases are used, then it may require row-level or even cell-level access control to support privacy constraints. An ambitious objective is to automatically generate fine-grained database access control and auditing policies from those in the middle tier, to eliminate potential logical errors during policy implementation. Furthermore, efficiency of policy evaluation and enforcement in the bottom tier is an important issue that must be addressed.

User side. The user-side components include user agents for preference specification and policy processing and presentation. The preference specification part interacts with the user through a paradigm that is close to the user's privacy protection objectives and generates privacy preferences in a formal language, so that the matching between enterprises' privacy policies and users' preferences can be conducted automatically. Additionally, the user agent provides an expressive user interaction model. When necessary, it presents the policies in an accurate and accessible manner and interacts with the user to help achieve privacy protection objectives.

Usability. This framework seeks to enable end users to take an active role in protecting their privacy online; thus, usability is a key component. Because maintaining security and privacy is heavily reliant on users' cooperation (that is, users need to specify their preferences), the maximal benefit of these preference specification methods cannot be realized unless interactions between the user and the system are simple and friendly. In particular, policy authoring and analysis tools as well as user agents need to be designed based on a comprehensive study of potential users' behaviors/preferences and existing tools.

Society, law, and economics. Social norms and laws serve as the fundamental guidelines for enterprises to regulate their privacy practices and for users to establish necessary information disclosure principles. Although many organizations now post online privacy policies, these organizations must realize that simply posting a privacy policy on their Web site does not guarantee compliance with existing legislation. To date, privacy protection law in the U.S. includes coverage for health care data (the Health Information Portability and Accountability Act, HIPAA), information obtained from and/or about children (the Children's Online Privacy Protection Act, COPPA) and financial data (the Gramm-Leach-Bliley Act, GLBA). They not only regulate the collection and use of private information inside one organization, but also concerns about cross-organization information sharing.

If privacy regulations and laws are systematically analyzed and mapped to formal semantics, common privacy practice pitfalls can be avoided. Additionally,

by studying users' social behaviors when accessing online services we will be better equipped to understand users' real privacy concerns—concerns they do not articulate, but are evident in their behaviors. Societal studies benefit individual users and enterprises because it helps them design user-acceptable privacy policies. Finally, economic factors play important roles to promote the consideration of privacy and the adoption of privacy protection technologies, especially in the enterprise side. There is a need to study enterprises' and users' behavior from the economic perspective.

RESEARCH ISSUES

Specification of privacy policies. Privacy policies in the top tier are contracts between enterprises and end users. A language for expressing such contracts must have unambiguous semantics and significant expressive power. As discussed here, existing specification languages for privacy policies lack both. Relevant research issues that need to be addressed in this context include the following:

Development of a formal language for specifying privacy policies. Although P3P's limitations have been widely acknowledged [10, 11], the exact limitations have not been clearly identified and no comprehensive solution has been proposed. A recent analysis of over 100 privacy policies in three different domains, for example, general e-commerce, health care, and financial Web sites [2], has yielded over 1,000 goal statements and identified the goals appearing most frequently in textual policies. Most of these goals cannot be expressed in current privacy languages and thus they can be used to drive the development of more expressive formal privacy languages.

It is also critical to develop expressive privacy policy languages with unambiguous semantics, serving as a semantic foundation for natural language privacy policies. An initial approach toward the definition of such semantics has been recently proposed [9], based on which integrity constraints are introduced to maintain a P3P policy's semantic consistency. That approach focuses on providing a formal semantics for P3P, rather than remedying other weaknesses of P3P. However, it is possible to build on that work in order to develop more expressive languages for specifying privacy policies, with a precise and clear relational semantics.

Automatic translation from natural language privacy policies to formal language policies. Although formal-language policies are being developed and deployed, it is unlikely they will replace natural-language privacy policies in the foreseeable future. Using existing natural language processing software, tools can be devel-

oped to translate natural privacy policies into formal-language policies. Such tools would also facilitate the automatic generation of formal-language policies to and from natural-language policies, and consistency checking between formal and informal policies as well as within natural-language policies, and would enable large-scale processing of online natural language privacy policies.

ENFORCEMENT AND AUDITING OF PRIVACY POLICIES

To guarantee an enterprise's systems are in compliance with its privacy policies in the top tier, privacy constraints must be integrated into specific applications in the middle tier to be effectively enforced in business operations. An enterprise often provides several services to its users, and information flow frequently happens between different applications. Thus, privacy policy enforcement and auditing should be considered not only in the context of a single application but also in the context of the information exchange between different applications/systems. The recent JetBlue Airways privacy breach further motivates this requirement [3].

Top-tier privacy policies are abstract and cannot be directly enforced in the middle tier. Thus, we must refine and materialize top-tier policies and map them into the relevant application domains. In particular, it is necessary to specify middle-tier privacy policies based on specific application models; verify their consistency with top-tier policies; and integrate middle-tier privacy policies with access control policies of underlying data management systems that ultimately control private information access. Relevant research issues that must be addressed in this context include the following:

Development of policy languages for specifying access control and auditing policies. Privacy protection requires either the design of new access control models or significant enhancement to current models. Most privacy policies allow users to decide whether to opt-out or opt-in to certain data usages; thus, a user's choices and consents must be stored and used to make access control decisions. As a result, the access requirement depends on both an enterprise's policies and user's choices. A language is needed to enable such highly fine-grained access control policies. The policy language should also specify auditing requirements for data access, so an audit trail can be generated. One research problem is the selection of an abstract data model. Another problem is the selection of a user model that allows access based on the attributes of users (for example, the roles the user is playing, the tasks the user is currently undertaking).

Theory and tools for comparing top-tier and middle-

tier policies. An enterprise needs to ensure that middle-tier policies correctly enforce high-level policies. One way to ensure this is to automatically generate middle-tier policies from high-level policies. However, because middle-tier policies contain more information than high-level policies, it is likely they cannot be automatically generated from high-level policies. It is likely that policies in the two tiers are specified separately; thus theory and tools for checking policy compliance must be developed. Such tools should also ensure that auditing policies in the middle tier are sufficient to generate an audit trail proving they are in compliance with high-level policies.

Algorithms and tools to automate translation from middle-tier to bottom-tier policies. Efficient enforcement of middle-tier policies requires the use of native access control and auditing mechanisms provided by the data storage program (databases). The Virtual Private Databases (VPD) feature in Oracle provides fine-grained access control as well as auditing by dynamically executing a policy, which is a PL/SQL program, and attaching the generated predicate to each query. While this allows very flexible policies, authoring policies involves writing complicated procedure programs—a highly error-prone process. Furthermore, it is difficult to verify whether the policies are implemented correctly. Therefore, a mechanism is required to automatically translate middle-tier policies into physical repository policies.

Theory for information flow control based on privacy policies. Different enterprise sectors often have different privacy policies in place. Such heterogeneity comes from several sources. Global enterprises may be subject to privacy laws from different countries. Company mergers may result in enterprises with distributed and heterogeneous information systems, which in turn may have heterogeneous privacy policies. However, because the various sectors of an enterprise are often interconnected, the information flows among these sectors must be properly controlled to prevent privacy breaches [3]. A key step in addressing this is the definition of a lattice based on privacy policies. This lattice definition will entail investigation of criteria and techniques for policy comparison. It is also important to investigate the extent to which the theory of information flow developed for MAC [5] can be applied. To actually deploy information flow control techniques, one must properly define the interacting entities in an information flow process. Such interacting entities can be defined in various ways—according to organizational functions or a technical point of view (that is, an entity can be an application program or a database system). Finally, a general notion of privacy contexts, which can be

We envision many cases in which human users would like to read the policy before entrusting their sensitive information to a Web site, rather than having a tool automatically make the decision for them.

defined as a component within an organization characterized by a homogenous privacy policy with respect to a given sets of data is needed.

PRIVACY MANAGEMENT FOR THE END USER

Privacy policies must be communicated to end users, enabling them to make meaningful decisions about whether to provide personal data online. However, just having the privacy policy in machine-readable form is only a first step toward enabling end users to control their privacy. We must develop a user interaction model and a user agent that interacts with the user through high-level objectives. Relevant research issues that must be addressed in this context include the following:

Development of a paradigm for specifying privacy preferences. This paradigm should be close to users' privacy objectives, rather than close to the data collection policies. Technical aspects of data collection and usage are often too complicated for users to fully comprehend. We conjecture that users' preferences should not be specified in terms of sharing specific data items, but rather in achieving privacy objectives. This paradigm should take into consideration users' limitations—it should be able to protect users from their own errors. The paradigm will account for privacy preferences that may vary for different transactions and Web sites. One possibility is to organize a set of preferences based on users' goals and Web sites' trust levels.

Methods and tools to present privacy policies to end users in a uniform and accessible way. The P3P effort is predicated on the belief that privacy policies are too difficult for humans to understand; thus they are encoded in machine-readable form, which is then automatically processed by tools. We envision many cases in which human users would like to read the policy before entrusting their sensitive information to a Web site, rather than having a tool automatically make the decision for them. Instead of presenting users with pages of text laden with legal terms and not

understandable to the majority of Internet users [2], privacy policies should be presented in summary form for the users. Once the most significant axes of users' privacy concerns and goals are determined, we can determine how to best structure, organize, and present this information to end users. For example, the presentation may include scenarios of what the company can do, warn of possible negative consequences, or stress differences with existing preferences.

PRIVACY POLICY: LEGAL AND ECONOMIC PERSPECTIVES

Existing privacy policies are largely driven by organizations' legal concerns. Moreover, different organization's policies address different issues, despite being in the same industry [2]. This suggests that companies within the same industry have different interpretations of the law or that errors of omission are common in privacy policies. In either case, while writing policies to address legal concerns is an understandable and prudent practice, it often leads to a mismatch between users' concerns and the information organizations disclose. Just as a law must survive constitutional challenge, a specified system should be demonstrably policy compliant. Part of the solution to helping financial institutions become GLBA compliant is for organizations to be able to show that policies meet the requirements of the law, and that they are complete and unambiguous [2].

CONCLUSION

Privacy is increasingly a major concern that prevents Internet users from fully enjoying the convenience, variety, and flexibility offered by online services. A variety of privacy-enhancing technologies has been proposed. While some technologies aim at preventing attacks that breach users' privacy, privacy policy technologies assume a cooperative relationship between service providers and users. Privacy policies allow enterprises and Internet users to communicate and negotiate privacy practices, and make online services privacy-aware. The proposed framework iden-

tifies key research challenges for the deployment and management of privacy policies. The framework shows that addressing these challenges will require close collaboration between academia and industrial researchers from multiple disciplines. **C**

REFERENCES

1. Acquisti, A. Privacy in electronic commerce and the economics of immediate gratification. In *Proceedings of the ACM Electronic Commerce Conference*. ACM Press, NY, 2004, 21–29.
2. Antón, A.I., Earp, J.B., Bolchini, D., He, Q., Jensen, C., and Stuffelbeam, W. The lack of clarity in financial privacy policies and the need for standardization. *IEEE Security & Privacy* 2, 2 (2004), 36–45.
3. Antón, A.I., He, Q., and Baumer, D. The complexity underlying Jet-Blue's privacy policy violations. *IEEE Security & Privacy* (to appear).
4. Ashley, P. A privacy logging and reporting framework. In *Proceedings of the Asia Pacific Information Technology Security Conference* (University of Queensland, May 24–27, 2004).
5. Bell, D. and LaPadula, L. Secure Computer Systems: Mathematical Foundations. *Technical Report MTR-2547* 1. MITRE Corporation, Mar. 1973.
6. Earp, J.B., Baumer, D.L., and Poindexter, J.C. Modeling privacy values with experimental economics. *Workshop on Privacy in the Electronic Society* (Washington, D.C., Oct. 28, 2004).
7. Karjoth, G., Schunter, M., and Van Herreweghe, E. Translating privacy practices into privacy promises—How to promise what you can keep. In *Proceedings of the 4th IEEE International Workshop on Policies for Distributed Systems and Networks*. June 2003, 135–146.
8. Karjoth, G., Schunter, M., and Waidner, M. Platform for enterprise privacy practices: Privacy-enabled management of customer data. In *Proceedings of the Second International Workshop on Privacy Enhancing Technologies*, LNCS 2482, (2003), 69–84.
9. Li, N., Yu, T., and Antón, A.I. A semantics-based approach to privacy languages. *CERIAS Technical Report TR 2003–28*. Purdue University, Nov. 2003.
10. Schunter, M., Van Herreweghen, E., and Waidner, M. Expressive privacy promises—How to improve the Platform for Privacy Preferences (P3P). Position paper for W3C Workshop on the Future of P3P (available at www.w3.org/2002/p3p-ws/pp/ibm-zuerich.pdf).
11. Schutzer, D.M. Citigroup P3P position paper. Position paper for W3C Workshop on the Future of P3P (available at www.w3.org/2002/p3p-ws/pp/ibm-zuerich.pdf).
12. W3C. Platform for Privacy Preferences (P3P) project; www.w3.org/P3P/

This work is funded by NSF Cyber Trust Grants #0430274 and #0430166.

ANNIE I. ANTÓN (aianton@ncsu.edu) is an associate professor of computer science at North Carolina State University, Raleigh, NC.

ELISA BERTINO (bertino@cerias.purdue.edu) is a professor of computer science and electrical and computer engineering at Purdue University, West Lafayette, IN.

NINGHUI LI (ninghui@cs.purdue.edu) is an assistant professor of computer science at Purdue University, West Lafayette, IN.

TING YU (yu@csc.ncsu.edu) is an assistant professor of computer science at North Carolina State University, Raleigh, NC.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.