

# Designing a Privacy Preference Specification Interface: A Case Study

Lorrie Faith Cranor

AT&T Labs-Research

180 Park Ave. Room A205

Florham Park, NJ 07932 USA

+1 973 360 8607

lorrie@acm.org

## ABSTRACT

User interface designers frequently face the challenge of making software configuration options accessible to users. In this paper I use my experience designing the interface for a Platform for Privacy Preferences (P3P) user agent as a case study to highlight some of the challenges interface designers face when developing configuration options for security and privacy software. In addition, I discuss several approaches to addressing these challenges.

## Keywords

P3P, privacy, user agent, preferences, software configuration

## INTRODUCTION

User interface designers frequently face the challenge of making software configuration options accessible to users. Organizing and presenting configuration options in such a way that users can find the options they need when they are most relevant can be difficult. Explaining to users the consequences of selecting each option can also be a challenge. Developing configuration interface for privacy and security software can be particularly problematic due to the fact that most users are not experts in these areas and are not even familiar with most of the terminology used to describe configuration options. In this paper I use my experience designing the interface for a Platform for Privacy Preferences (P3P) user agent as a case study to highlight some of the challenges interface designers face when developing configuration options for security and privacy software.

## THE PLATFORM FOR PRIVACY PREFERENCES

The Platform for Privacy Preferences (P3P) is a World Wide Web Consortium (W3C) recommendation that provides a standard, computer-readable, XML format in which web sites can express their privacy policies as well as standard mechanisms for web browsers and other P3P “user agent” software to locate and fetch these policies

[4,5]. P3P user agents may check for P3P policies at web sites a user visits, compare them with users’ previously specified privacy preferences, and provide feedback to the user about these policies. Some user agents make cookie-blocking decisions on the basis of this comparison or take other actions such as allowing or denying access to a user’s electronic wallet.

The P3P 1.0 Specification [5] defines a P3P “vocabulary” that includes eight major components, most of which contain multiple sub-components and attributes. Each component is represented as an XML element. For example, the use of collected data is represented by the “PURPOSE” element. The specification defines 11 purpose sub-elements, each representing a different data use. In addition, each of these purpose sub-elements has a “required” attribute that indicates whether the data may be used for this purpose all the time, on an opt-in basis, or on an opt-out basis.

P3P policies were designed both to provide information about website privacy policies that a human might use to make decisions (such as whether or not to shop at a particular web site or whether to exercise “opt-out” options), and to facilitate automated decision-making (such as whether to display a privacy warning or whether to block cookies at a particular web site). The details of how a P3P user agent might use a P3P policy to display information or to make automated decisions are not part of the P3P specification. User agent implementers thus face questions about how much information to present, what words and phrases to use, what aspects of privacy policies to emphasize, and how to make this information most accessible to end users. They also face questions about how to elicit privacy preferences from users, the range of configuration options to offer, and the types of decisions that should be automated. These questions can be grouped into two major interface design challenges: an interface for informing users about web site privacy policies, and an interface for configuring a P3P user agent to take actions on the basis of a user’s privacy preferences. While our research has investigated both of these areas, the latter challenge is the focus of this paper.

## Privacy Bird

We developed the AT&T Privacy Bird [1] P3P user agent as a browser helper object [7] for the Internet Explorer 5.01, 5.5, and 6.0 web browsers on Microsoft Windows 98/2000/ME/NT/XP operating systems. We distributed the beta 1.1 version as a 1.4 MB self-extracting file that includes an installation wizard. Once installed, a bird icon appears in the top, right-hand corner of the user's Internet Explorer browser windows. The bird icon changes shape and color to indicate whether a web site is P3P-enabled, and (if it is P3P enabled) whether its privacy policy matches a user's privacy preferences. Users can click on the bird to access additional information about the current web site's privacy policy as well as configuration and help menus. When a user selects the privacy configuration menu item a privacy preference specification interface appears, as shown in Figure 1.

## DESIGN CHALLENGES

Designing a user interface for specifying privacy preferences is difficult for several reasons: privacy policies are complex, user privacy preferences are often complex and nuanced [3], users tend to have little experience articulating their privacy preferences, and users are generally unfamiliar with much of the terminology used by privacy experts.

## Complex Privacy Policies

P3P privacy policies include eight major components, most of which include sub-components. Some components are represented as elements for which there are fixed sets of possible values, while other components are represented by elements that can include text strings or extensible sets of possible values. User privacy preferences often reflect a combination of privacy policy components. For example, a user may wish to receive a privacy warning at sites that collect financial information and use it for marketing, but not at sites that collect financial information and use it to process an order nor at sites that collect preference information and use it for marketing. Even if we limit our discussion to those elements with fixed sets of possible values and ignore the attributes that may modify these elements, there are over 36,000 possible combinations of privacy policy components that can be expressed using the P3P syntax. Potentially, users may wish to express a preference over any of these combinations.

## Complex Privacy Preferences

Surveys have repeatedly shown that the majority of people take a pragmatic approach to privacy, making contextual decisions about whether to protect their privacy or take actions that might put their privacy at risk [3,9,10]. Furthermore, empirical studies have found that Internet users' behavior is often inconsistent with their self-reported privacy preferences [14]. This suggests that, in practice, users are willing to make privacy tradeoffs that may be difficult for them to specify in advance. For example, users may have a preference not to have their web browsing activities monitored and used to build profiles about them. They may feel particularly strongly about this when visiting health and medical web sites, but they may be willing to allow this monitoring at web sites of book and music retailers that use this information to make personalized recommendations of books and CDs and offer discounts. Indeed some users who otherwise eschew monitoring may even request such monitoring if they find the recommendation service particularly useful.

## Inexperienced Users

While most people will readily proclaim a desire for privacy, they usually have little articulating a comprehensive set of privacy preferences or rules for a user

Figure 1. AT&T Privacy Bird privacy preference specification panel

agent. The task is difficult even when limited to specifying privacy preferences with respect to web site interactions (the only concern of most P3P user agents). Furthermore, most Internet users have little understanding about how online profiling is done, how cookies work, or what the real online privacy risks actually are. Thus they may be ill equipped to create detailed specifications of privacy preferences.

The task of specifying privacy preferences is further complicated by the fact that discussions of privacy often involve jargon. Many privacy policies are written in language that is understandable only to privacy experts and lawyers. Readability experts have found that the privacy policies on many popular web sites are written at a college reading level or higher [13]. Thus, it is quite understandable that Internet users report finding privacy policies difficult and time consuming to read [12]. Interface designers are challenged with designing a privacy preference specification interface that uses readily understandable language and avoids the jargon commonly found in privacy policies.

### **APPROACH TO DESIGN CHALLENGES**

In order to design a privacy preference specification interface useful to users with little experience specifying complex privacy preferences we had to find ways of reducing the complexity and focusing in on the issues that users would likely be most concerned about. We focused our efforts on a subset of the P3P vocabulary, bundled together similar vocabulary elements with distinctions unlikely to be important to users, and used terminology free of jargon. We also created privacy options that used combinations of P3P data elements, and used layering to allow experts to provide additional privacy options. Our approach allowed us to minimize the use of defaults.

#### **Vocabulary Subset**

We made the P3P vocabulary appear simpler by designing an interface that highlights a subset of the P3P vocabulary that is likely to be of most interest to users. We reduced the number of options by eliminating those combinations of options unlikely to be useful in practice. We focused on the data practices that seemed to raise the most concerns for American Internet users: collection of health and financial information, marketing, profiling, and sharing personal data with other companies. Although location information is also very sensitive, we did not include a setting that dealt with that information in the AT&T Privacy Bird interface because we do not anticipate that this user agent will be used on devices where location information is an issue. It might be important to highlight other data practices in user agents designed for other types of users (children, non-Americans, etc.) or for mobile devices.

#### **Bundling Similar Vocabulary Elements**

Many of the distinctions made in the P3P vocabulary are unlikely to be important to most users—although it is quite likely that the distinctions users find most important will change over time and perhaps even vary across regions of the world. We bundled vocabulary elements together that users may think about in similar ways in order to reduce

the apparent complexity of the P3P vocabulary. For example, we bundled the six types of P3P data recipients into two groups—sharing, and non-sharing—and described the sharing practice as sharing data “with other companies (other than those helping the web site provide services to me).” Sites that disclose data only to their agents and to delivery companies are considered to be non-sharing, while those that disclose data to any other recipients are considered to be sharing. Thus P3P vocabulary distinctions between sites that share data with companies having similar privacy policies, companies having different privacy policies, and companies with unknown privacy policies are hidden from Privacy Bird users. For P3P experts who want to understand how exactly our bundles map onto the P3P vocabulary, we provide detailed information in the accompanying help files.

#### **Removing Jargon**

The P3P vocabulary terms borrow terminology from privacy laws and fair information practice principles. While these terms are well known to privacy experts, they are foreign to almost everyone else. Thus it is a challenge for user agent implementers to come up with terms that will be more meaningful to users, while accurately describing the P3P vocabulary.

The P3P vocabulary also uses terms such as “pseudonymous analysis” and “individual decision,” which are meaningless without their accompanying definitions, even to privacy experts. These definitions are too lengthy to be used verbatim in a user interface. We experimented with approaches to describing these purposes that privacy advocates consider to be variations on “profiling.” However, the term “profiling” did not appear to be any more meaningful to most users than the vocabulary terms. From a privacy perspective, it is very important to know that these purposes involve building a record about an individual. However, a description of what the record might be used for seemed to resonate better with users. Ultimately we ended up bundling the profiling purposes with the marketing purposes and some of the most sensitive data groups and the setting became “Warn me at web sites that use my [data category] information for analysis, marketing, or to make decisions that may affect what content or ads I see, etc.”

#### **Using Vocabulary Elements in Combination**

Internet users tend to have complex privacy preferences that generally cannot be captured by focusing on a single dimension of the P3P vocabulary. It is therefore important that privacy preference options reflect this complexity. For example, we limit warnings about the collection of health and medical information to sites that use this information for purposes that we believe users will most likely find objectionable (marketing, profiling, and sharing with other companies). As a result users should not get warnings at health web sites unless those sites collect health data for one of these objectionable purposes. Indeed eight of our 12 warnings are triggered by a combination of data practices rather than the presence of a single P3P element.

## Layered Interfaces

A common way of reducing the complexity of software user interfaces is to divide the interface into two or more layers. Many pieces of software feature configuration menus that include only the most commonly used settings, and a separate “advanced” menu that includes the less frequently used settings. This is an effective way to hide complicated options from users who will never need to access them; however, it sometimes becomes difficult for users who want to access advanced settings to find what they are looking for. In the Privacy Bird interface we show a choice of three “pre-packaged” privacy settings on the configuration screen, in addition to 12 check boxes that users may use to select “custom” settings. When a user selects one of the pre-packaged settings, the boxes next to the corresponding custom settings are checked automatically. This provides immediate feedback about what each of the pre-packaged settings does. In addition, it makes it easy for users to make modifications to a pre-packaged setting.

The ability to import privacy settings using APPEL [6] adds another layer to P3P user agents. The APPEL language allows for much more detailed configuration options than most graphical user interfaces can support.

## Default Settings

Despite our efforts to develop usable configuration interfaces, most users rarely change the default settings on many of the software packages they use. Changing the settings can be time consuming and confusing [11], and users risk “messaging up” their settings and being unable to return their software to the state they have grown accustomed to. Designers face choices not only about what the default settings should be, but also when to employ defaults and when to “force” users to make choices [0].

In our design, we tried to avoid setting defaults for the main privacy settings because we wanted users to select settings that would reflect their personal privacy preferences. We wanted to force users to choose the settings themselves; however, we were concerned that it would be difficult for users to make such choices before they had spent time using and understanding the software. So we decided to offer users only the high, medium, and low options during software installation, and make all of the custom options available after the software was installed. However, users complained that they wanted more information about these settings during the installation process, so we plan to provide the full configuration options in our next release.

## CONCLUSIONS

In this paper I have reviewed some of the design challenges faced while developing the AT&T Privacy Bird privacy preference specification interface. I have highlighted several design approaches that we found useful and believe may have utility in future privacy preference specification and security configuration interface development efforts.

## REFERENCES

1. Cranor, L. Arjula, M., and Guduru, P. Use of a P3P User Agent by Early Adopters. *Proceedings of the*

ACM Workshop on Privacy in the Electronic Society, (Washington, DC, November 2002), ACM Press.

2. Cranor, L. and Reagle, J. Designing a Social Protocol: Lessons Learned from the Platform for Privacy Preferences Project, in J.K. MacKie-Mason and D. Waterman (eds.) *Telephony, the Internet, and the Media*. Lawrence Erlbaum Associates, Mahwah, NJ, 1998. Available at <http://www.w3.org/People/Reagle/papers/tprc97/tprc-f2m3.html>
3. Ackerman M.S., Cranor, L.F., and Reagle, J. Privacy in E-Commerce: Examining User Scenarios and Privacy Preferences, in *Proceedings of EC'99* (Denver CO, November 1999), ACM Press, 1-8.
4. Cranor, L. *Web Privacy with P3P*. O'Reilly & Associates, Sebastopol CA, 2002.
5. Cranor, L., Langheinrich, M., Marchiori, M., Presler-Marshall, M., and Reagle, J. *The Platform for Privacy Preferences 1.0 (P3P1.0) Specification*. World Wide Web Consortium Recommendation, April 2002. Available at <http://www.w3.org/TR/P3P/>.
6. Cranor, L., Langheinrich, M., and Marchiori, M. *A P3P Preference Exchange Language 1.0 (APPEL1.0)*. World Wide Web Consortium Working Draft, April 2002. Available at <http://www.w3.org/TR/WD-P3P-Preferences>.
7. Esposito, D. Browser Helper Objects: The Browser the Way You Want It, MSDN Library, January 1999. Available at <http://msdn.microsoft.com/library/default.asp?url=/library/en-us/dnwebgen/html/bho.asp>.
8. Georgia Tech Graphics, Visualization & Usability Center. GVU's 10th WWW User Survey, 1998. Available at [http://www.gvu.gatech.edu/user\\_surveys](http://www.gvu.gatech.edu/user_surveys)
9. Harris, Louis and Associates and Westin, A.F. Harris-Equifax Consumer Privacy Survey 1991. Equifax Inc., Atlanta GA, 1991.
10. Harris, Louis and Associates and Westin, A.F. E-commerce & Privacy: What Net Users Want. Privacy & American Business, Hackensack NJ, 1998.
11. Mackay, W.E. Triggers and barriers to customizing software, in *Proceedings of CHI'91*, ACM Press, 153-160.
12. Privacy Leadership Initiative. Privacy Notices Research Final Results. Conducted by Harris Interactive, December 2001. Available at <http://www.ftc.gov/bcp/workshops/glb/supporting/harris%20results.pdf>.
13. Rodger, W. Privacy Isn't Public Knowledge: Online policies spread confusion with legal jargon, *USA Today*, 1 May 2003, 3D. Available at <http://www.usatoday.com/life/cyber/tech/cth818.htm>.
14. Spiekermann, S., Grossklags, J., and Berendt, B. E-Privacy in 2<sup>nd</sup> Generation E-Commerce: Privacy Preferences versus actual Behavior, in *Proceedings of EC'01* (Tampa FL, October 2001), ACM Press, 38-47.