RC23495 (W0501-070) January 18, 2005
Computer Science

IBM Research Report

Designing a Privacy Management Tool: Progressive Human-Computer Interaction Based Privacy Research with Organizational Users

Carolyn Brodie, John Karat, Clare-Marie N. Karat
IBM Research Division
Thomas J. Watson Research Center
P.O. Box 704
Yorktown Heights, NY 10598

Research Division
Almaden - Austin - Beijing - Haifa - India - T. J. Watson - Tokyo - Zurich

Designing a Privacy Management Tool: Progressive Human-Computer Interaction Based Privacy Research with Organizational Users

Carolyn Brodie, John Karat, and Clare-Marie Karat
brodiec, jkarat, and ckarat@us.ibm.com
IBM T. J. Watson Research Center
Hawthorne, NY 10532

November 15, 2004

ABSTRACT

Usability has been identified as a major challenge to moving the results of security and privacy research to use in real systems [15]. One reason seems to be that there has been only limited research into how to make complex security and privacy functionality understandable to those who must use it. The research reported here describes our efforts to design a system which facilitates privacy policy authoring, implementation, and compliance monitoring. We employed a variety of user-centered design methods with 109 users across the five steps of the research reported here. The majority of these users are organizational privacy professionals. This case study highlights the work of identifying organizational privacy requirements, iteratively designing and validating a prototype with users, and conducting laboratory tests to guide specific design decisions for flexible privacy enabling technologies. Each of the five steps in our work is identified and described, with a particular emphasis on the motivation for each step and the user-centered methods employed. Recommendations for extending this work into the security arena are included.

INTRODUCTION

The rapid advancement of the use of information technology in industry, government, and academia makes it much easier to collect, transfer, and store personal and business information around the world. At the same time threats to information used, stored, and shared within and between organizations and individuals have been steadily growing. While an increasing amount of research concentrates on identifying security and privacy weaknesses and how to address them, making this technology usable remains an issue. The Computing Research Association (CRA) Conference on Grand Research Challenges in Information Security and Assurance has identified the ability to 'give end-users security controls they can understand and privacy they can control for the dynamic, pervasive computing environments of the future' as a major research challenge [15]. As Whitten and Tygar [39] point out, 'security mechanisms are only effective when used correctly' and that this is often not the case due to usability issues with security software. In this paper we highlight how human computer interaction (HCI) research can help to address these weaknesses through a case study of the design of a set of utilities for privacy policy management. In particular, we discuss the creation of a set of privacy utilities that is designed to assist organizations with the creation, implementation, and internal auditing of privacy policies. During this project user-centered design methods have been employed from the beginning to ensure that the functionality and interaction methods exposed to the user match the users' skill set and understanding of the task.

We chose the domain of organizational privacy policy creation and enforcement because use and misuse of personal information (PI) is an area of increasing concern in many geographies and domains around the world. Organizations depend on the use of PI to meet the needs of their customers, constituents, and patients in an efficient manner. This raises challenging questions and problems regarding the use and protection of PI [25]. Questions of who has what rights to information about us for what purposes become more important as we move toward a world in which it is technically possible to know just about anything about just about anyone. As stated by Adams and Sasse [1]: 'Most invasions of privacy are not intentional but due to designers' inability to anticipate how this data could be used, by whom, and how this might affect users.' When designing systems that use personal information, we must not only secure them so that information cannot be accessed by unauthorized users but also from authorized users for unauthorized purposes.

We provide a description of our scope in addressing this area, since privacy can and does mean different things to different people. We are primarily focused on a view of privacy as the right of an individual to control personal information use rather than as the right to individual isolation [28, 30, 36]. Organizations commonly provide a description of what kind of information they will collect and how they will use it in privacy policies. In some areas (e.g., the collection and use of health care information in the US or movement of personal information across national boundaries in Europe) such policies can be required, though the content of the policy is not generally specified in legislation. While there has been considerable consensus around a set of high level privacy principles for information technology [30], we do not think it is likely that a single privacy policy can be created to address all information privacy needs. For example, there will likely be considerable differences in privacy legislation in different regions of the world [26]. Similarly, organizations in different fields (e.g., health care, banking, government) need to tailor policies to their domains and needs [10, 11].

In this paper we will present two years of research in which we employed a variety of user-centered design methods with 109 users in order to understand and meet the needs of organizational users tasked with the creation, implementation, and auditing of privacy policies. The majority of these users are organizational privacy professionals. We will report how we identified organizational privacy requirements, analyzed privacy architectures and identified missing elements needed by organizational users to effectively create, implement and audit privacy policies, iteratively designed and validated a prototype with users, and conducted laboratory tests to guide specific design decisions. Finally, we will discuss recommendations that we have developed for extending this work into the security software arena.

RELATED WORK

There are many aspects of privacy that have been the subject of research, including research on the public perceptions of the need to protect PI, research and development of many types of privacy preserving technologies, as well as research into the current approaches that are being used by organizations to protect the PI of their customers, constituents, patients, and employees. In this section we will discuss recent research into the public perceptions of privacy within organizations and how they affect individual willingness to share data, technological approaches for enforcing privacy policies, and finally how organizations are protecting PI today.

Research has identified high levels of consumer concerns regarding privacy [17, 18, 31]. A multi-national consumer privacy survey in 1999 investigated US, German, and UK consumers' attitudes toward privacy in different industries [18]. Although customer confidence levels varied across different industries, generally consumers expressed fears about the misuse of their PI. Seventy-eight percent of the people in the survey reported that they have refused to provide information in the past due to concerns about PI misuse. A privacy and business survey in 2000 conducted for the Australian government revealed that 95% of the respondents think it is necessary to implement laws to protect PI and also documented that approximately 50% of the respondents routinely and intentionally provide inaccurate PI [31]. A more recent Forrester report found that 97% of North American consumers believe that online privacy concerns are real and 94% reported that they believe the benefits they receive for sharing personal information do not outweigh their concerns [14]. In the health care domain, physicians and practitioners are concerned about serious threats to patient privacy due to information gathering methods, record accuracy and access, and unauthorized secondary use [11]. In the education sector, a Stanford University report reveals that PI is not effectively protected [36].

Researchers have responded to these concerns through the development and analysis of machine readable privacy policies and the development of mechanisms for helping end-users to understand the policies and organizations to enforce the policies. One area of research is on the development and use of machine readable privacy policy schemas for enabling privacy functionality. P3P [16] is one of the first privacy policy languages that has been standardized by an international standards body, the W3C. P3P is an XML based language that allows organizations with Websites to create machine readable versions of their privacy policies. Generally, P3P allows organizations to specify rules that contain the type of data, the type of use, the user of the data, the purpose of the use, and how long the data will be retained. From the end-user or client point of view, automated agents, such as the AT&T Privacy Bird [8] and browsers such as Microsoft's Internet Explorer [27] can use the P3P policies to provide individual users with the ability to quickly determine if the Website's privacy policies match their privacy preferences. Other proposed schemas, such as APPEL, have expanded on the goal of helping individuals to quickly determine if a Website's policies match their preferences by allowing the user to define rule sets for describing acceptable organizational privacy policies [40].

While the ability to quickly understand a site's privacy policy and determine if the site conforms to their preferences is helpful to end-users, it is important to understand that there is no guarantee that the policy is actually implemented as specified within the organization. This fact has lead to research into how machine readable (XML schema languages) privacy policies can be used by organizations to enforce policies. Karjoth and Schunter [24] analyzed how enterprise privacy policies differ from security policies and how well P3P can express an enterprise privacy policy. Based on this analysis, they propose a privacy policy model that can be used for internal access control within an enterprise. New XML schemas designed to enforce privacy policies include, the Enterprise Privacy Authorization Language (EPAL) [7] and XACML with a privacy profile [29]. These allow more expressive policies that include hierarchical policy elements, conditions on rules, and a user definable set of obligations. EPAL is being considered by the W3C standards body and XACML with a privacy policy profile is being considered by OASIS. The ability to use a language like EPAL to capture and logically enforce the privacy policies of large, complex organizations has been studied and formalized by Backes, Pfitzmann and Schunter [9].

In addition to policy analysis, researchers have been exploring enforcement mechanisms for some time. Anderson [4,5] proposed a security policy model for the British Medical Association that described how to implement and manage compartmented security in health care. In an update in 2000, he reported that it had been implemented successfully in three British Hospitals [5]. Since that time there has been research into how machine readable policies can be used internally by organizations to enforce their privacy policies. Some approaches have concentrated on allowing policies defined by individuals to dictate how their information is used [13], while many others have concentrated on enforcing privacy policies created at the organizational level. An example of this is the Hippocratic Database [3] in which P3P is used to define access rules that are then enforced by the Hippocratic Database. IBM's Tivoli Privacy Manager is another example of an approach that has used P3P to define privacy policies which are then enforced by deploying monitoring software around data stores that sends requests for PI to a server which then determines if the access conforms to the privacy policy and logs both the attempt and the enforcement decision [19].

Even with all of the research that indicates that there is growing concern about privacy issues and the possible technical approaches that have been developed to protect PI, most organizations that depend on the use of personal information in their business processes have done little to implement the policies through technology [22, 35]. Privacy policy enforcement remains largely a human process. According to a 2003 study conducted by Ponemon for the IAPP [32] only 19% of the organizations sampled report that they are currently using any privacy enabling technology. This confirms the situation described by Forrester with respect to privacy [17]. This Forrester report reveals a mismatch between consumer demands for privacy and enterprise practices in industry. According to this report, although customer concerns about privacy remain high, the majority of executives (58%) believe privacy issues are addressed extremely well by their companies. Most executives don't know whether their customers check the privacy policies or not and few see the need to enhance their privacy practices. Research in the Asia-Pacific region complements these results [31].

The reality is that many organizations recognize that privacy is an issue for them. They currently do not know how to use technology to help them enforce their privacy policies. The Ponemon study [32] reported that although 98% of the companies in their survey have a privacy policy, 52% believe they do not have the resources to adequately protect privacy. Furthermore, most organizations store PI in heterogeneous server system environments and currently they do not have a unified way of defining or implementing privacy policies that encompass data collected and used by both Web and legacy applications across different server platforms [6]. This makes it difficult for the organizations to put in place proper management and control of PI, for the data users to access and work with the PI inline with the privacy policies, and for the data subjects to understand rights regarding use of their PI. It has been suggested that one reason that organizations are not employing new privacy enabling technologies to protect PI is that these technologies are currently very difficult to use [15,39]. In practice user-centered design techniques have contributed to the development of some highly usable security systems [21, 41]. Based on this evidence, our emerging focus has been on applying HCI-based research techniques to answering how organizations could create a wide range of policies, and how technology might enable the policies to be enforced and audited for compliance. We have elected to focus on technology to enable usable privacy policy authoring and enforcement, rather than trying to directly address what privacy rights people should have [e.g., 37] or how to de-identify information stored in systems [e.g., 33]. This does not mean that we think privacy is not an important social issue. Rather it points to our belief that technology can enable flexible, reliable and accountable privacy policies (i.e., be privacy enabling) and not just be a force which reduces individual rights. We hope our work contributes positively to this goal.

USER-CENTERED PRIVACY RESEARCH PROJECT

With the recognition that both individuals who share data and organizations who use data have concerns about how to protect PI and that many potentially valuable privacy enabling technologies are being researched and developed, but are not being widely used, we have executed a user-centered design research program on organizational privacy capabilities. Our research is rooted in approaches developed in the HCI field over the past 20 years to inform the design of usable systems [12, 34]. We employed a variety of usability methods to progress from identifying organizational privacy concerns and needs to designing and evaluating prototypes and design trade-offs. This work included (1) identifying privacy needs within organizations through email survey questionnaires, (2) refining the needs through in-depth interviews with privacy-responsible individuals in organizations, (3) studying proposed privacy technologies and architectures to determine how well these solutions meet the organizational needs we identified and hypothesizing about ways to address the gaps we found, (4) designing and validating a prototype of a technology approach to meeting organizational privacy needs through onsite scenario-based walkthroughs with users, and (5) collecting empirical data in a controlled usability laboratory test to understand the usability of privacy policy authoring methods included in our proposed design. These activities were completed between the early Spring of 2003 and Fall of 2004 and involved participation of 109 users.

Step One: Identifying the Needs

The initial interview research was completed in two steps: 1) an email survey of 51 participants to identify key privacy concerns and technology needs, and 2) in-depth interviews with a sub-set of 13 participants from the original sample to understand their top privacy concerns and technology needs in the context of scenarios about the flow of PI through business processes used by their organizations. For the first step, the research team recruited 51 participants, including 23 from Industry and 28 from Government, from North America, Europe, and Asia who were identified as part of organizations that represent innovators concerning privacy who could assist in the early identification of issues and requirements. The domains represented by the participants ranged from banking/finance, to education, health care, information technology, insurance, mass media, and travel/entertainment. Many of the participants from Industry represented multi-national companies that are in the process of understanding how to comply with privacy legislation around the world. The respondents from Government represented career public service employees from local, mid-level and central or national governmental organizations. They were recruited through a variety of mechanisms including follow-ups on attendance at privacy breakout sessions at professional conferences and referrals from members of the professional privacy community.

The email survey included three questions on privacy. They were as follows: 1) What are your top privacy concerns regarding your organization, 2) What types of privacy functionality you would like to have available to address your privacy concerns regarding your organization, and 3) At this time, what action is your organization preparing to take to address the top privacy concerns you listed above. Participants were asked to rank order their top three choices on questions 1 and 2 from a list of choices provided. They were also allowed to fill in and rank an "Other" choice if their concern was not represented in the list. For question 1 options included types of data they needed to protect, types of users seeking data access, types of legislation they were concerned with, and types of risk they felt concerned about. On the second question, options included different technology approaches to privacy management. On question 3, they were asked to select all that applied from a list of options which included internal development, external purchase, and external consulting activities, and again, they could select "Other" and describe the action they were taking if it was not listed as a choice. The data were grouped into Industry and Government segments. All participants in this and all other steps were promised that their data would be kept confidential and only used in a summarized or de-identified format.

For the first question, there was a statistically significant difference in the distributions of the top concerns of Industry and Government respondents [$X^2_{(10)}$=20.6, p<=0.025]. For the industry respondents, the top privacy concern involved "The economic harm that would result to this company if a privacy breach regarding customer data became public". In contrast, Government respondents stated that "Keeping external users from violating the privacy of others' data" was their top concern. Industry and Government shared two of the top three concerns.

These were “Keeping internal employees from violating the privacy of other’s data” and “Protecting the privacy of legacy data from unauthorized review or use”. For most business organizations, their brand is priceless. Thus possible economic harm to their businesses due to the adverse publicity that privacy breaches generate was their top concern. Government respondents did not see economic harm as a top privacy risk. Industry and Government respondents were concerned about the misuse of PI by their own employees and from external users.

The second question addressed the privacy functionality desired by respondents to address their top privacy concerns. There were no statistically significant differences in the distribution of responses on this question between Industry and Government. Industry and Government shared the four top-ranked choices, although they were ordered differently. The top-ranked privacy functionality included “One integrated solution for legacy and Web data”, “Application-specific privacy policy authoring, implementation, auditing, and enforcement”, “The ability to associate privacy policy information with individual data elements in a customer’s file”, and “Privacy protection for data stored on servers from IT staff with no need to view data content”.

The final question in the survey asked respondents about the actions being taken by their organizations regarding privacy. Again, no statistically significant differences by segment were found. The most frequent response was “We have begun to implement privacy solutions, but more work is needed”. Other typical responses included “We are going to bring privacy consultants in to advise our organization on how to proceed”, “We will develop a privacy solution in-house to address these concerns”, and “We will purchase a privacy solution to address these concerns”. In order to better understand participants’ views of privacy and to ground the interview data in the context of organizational use of PI in their business processes, we conducted in-depth interviews with a sub-sample of the respondents in Step Two.

Step Two: In-Depth Interview Research

The goals of the in-depth interviews were to build a deeper understanding of the participants’ and their organizations’ views regarding privacy, their privacy concerns, and the value they perceived in the desired privacy technology they spoke of in the context of scenarios of use involving PI in their organizations. The majority of the interview sessions were centered on discussion of a scenario of use provided by the respondent regarding PI information flow in their organization and follow-up questions related to it. We wanted to identify and understand examples of how PI flowed through business processes in the organization, the strengths and weaknesses of these processes involving PI, the manual and automated processes to address privacy, and the additional privacy functionality they need in the context of these scenarios. The participants whom we interviewed came from the finance/banking (4 interviewees), government (3), health care (2), travel/entertainment (3), and information technology (1) domains. The thirteen participants represented privacy officers, business process owners, information technology (IT), security, and privacy leaders. Most of the organizations represented were multi-national companies. Three of the thirteen individuals were part of organizations that were headquartered in Canada or Europe. The remaining participants represented multi-national businesses or government organizations headquartered in the US. All of the interviews were completed by telephone in about an hour. The interviews consisted of seven open-ended questions and resulting follow-up questions.

In their scenarios about the flow of PI through the organizational business processes, the interviewees shared a view that privacy protection of PI depends on the people, the business processes, and the technology to support them. In regard to the people, interviewees identified the need for education of data users about organizational privacy policies and practices that they would be responsible for upholding day-to-day in their jobs. Interviewees also highlighted the need to inform the data subjects of the content of the privacy policies and believed that the trust of the data subjects is critical and can be strengthened through this education and experience of privacy in practice. Interviewees believed that building trust with their customers provides a competitive advantage and is essential to their customers’ willingness to use personalization. This result confirms previous research [23].

Interviewees noted that their business processes are in need of redesign to implement the privacy policies within them. As an initial step, it is necessary to create data flows of the PI through the business processes and develop manual procedures to make the privacy policies operational. Interviewees noted three benefits in doing this redesign work. First, they have identified methods for improving their security as they work on imbedding privacy in the processes. Second, they have identified and removed redundancies in the collection of PI and state that this streamlining effort has saved them money and improved their relationships with data subjects, as one carefully protected view of each data subject has emerged. Third, the PI the organization has is more accurate and the organizations view it as a valuable asset and for businesses, which provides a competitive edge.

As the interviewees look to develop or acquire technology to automate the creation, enforcement and auditing of privacy policies, they have noted a number of concerns. They need to understand what PI data is located where in their organization, and complete data classifications of it. Organizations need strategies and business analyses regarding the risks involved in data consolidation and possible privacy breaches that may result from it. They need to protect PI from IT staff who have no need for database field level access of PI data. Finally, most applications and systems do not have an easy way to define retention dates for information. Interviewees are currently using a mix of manual and automated security functionality (tweaked to do things it was not originally intended to do) to implement their privacy policies. There are large gaps, particularly concerning the ability to perform privacy compliance checking.

The authors developed a set of five key privacy concepts based on the interviews and further research that are needed to meet the need of organizational users of privacy protecting technologies. They include:

1. It is important to provide users with **one integrated solution for an organization's heterogeneous configuration** even if it consists of a set of utilities that provide users with a similar set of functionality and interaction methods for systems that are implemented differently on different technologies.
2. The **privacy functionality must be separated from the application code** for cost, consistency, and flexibility – users do not want to have to modify all of their applications individually to ensure that PI is protected.
3. There needs to be the ability to **support an appropriate granularity** for applying the privacy policy (field level in a database)
4. There must be the ability to work with **structured and unstructured information**, protecting field level data and handling PI within documents in appropriate ways.
5. There must be **simple and flexible privacy functionality** that is designed to meet the needs of the user community that owns each subtask in the privacy process. For example, CPO's and/or business process owners often write the privacy policies. They must be able to author policies that will end up in machine readable form without having IT skills.

In Step Three, we will discuss how using these key privacy design concepts developed in the first two steps the authors identified areas within an abstract privacy architecture where user-centered design techniques can more effectively facilitate the design and development of a set of privacy utilities for allowing organizations to capitalize on the rich and diverse range of privacy technologies that are being researched and developed.

Step Three: Architectural Analysis of Privacy Functionality

Using the wealth of data collected from the survey and interview research, we identified a set of key design concepts for any privacy solution and used them to analyze existing privacy architectures to identify areas in which user-centered design techniques could be applied to best meet the needs of organizational privacy users. A key goal in any user-centered research project is to understand if and how well existing and new technologies meet the needs of the user community. In this step we analyzed possible privacy architectures to determine how well they meet the needs of organizational users concerned with protecting PI. To facilitate the description of this analysis we have created a generalization of many approaches to protecting the privacy of PI which is shown in Figure 1. In this figure a privacy policy authoring utility is used to create privacy policies that are stored in a machine readable format. This machine readable privacy policy is then used by a privacy enforcement mechanism that is positioned between applications and data stored within the organization's configuration. The architecture also provides for creation of a log of privacy events by the enforcement mechanism, which can be analyzed by the organization's audit mechanism to report on compliance with privacy policy. The generalized architecture drawing in Figure 1 is purposefully abstract so that it can be used to describe the common elements and mechanisms in a variety of possible privacy implementation approaches.

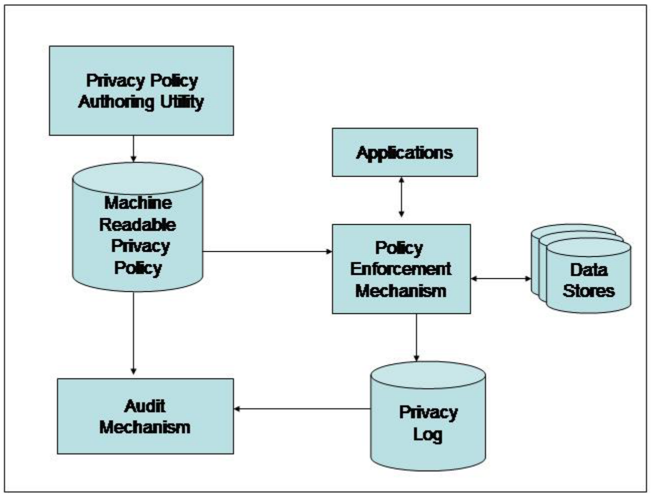
Privacy Policy Authoring Utility
Applications
Machine Readable Privacy Policy
Policy Enforcement Mechanism
Data Stores
Audit Mechanism
Privacy Log

Figure 1. Abstract Privacy Architecture

Different types of machine readable policies have been proposed and are at different points in the standardization process. P3P is currently a privacy language standard and is used to define privacy policies in some approaches [3, 19]. Other standards that allow for more expressive policies, such as EPAL and XACML with a privacy profile are also being considered. Likewise, there are many different approaches to privacy policy enforcement that have been proposed including query re-writing [2], data access monitoring and the use of a rules based enforcement engine [19], and the application of a modified access control mechanism [4,5]. Also, not all approaches in the literature include all components in the drawings. For example, the model proposed by Anderson uses an enforcement mechanism based on the concepts for multilevel security research as well as an audit mechanism, but does not address the use of machine readable privacy languages. While we recognize that each of these types of solutions do have the potential to be valuable to organizational users functionality, we have found that all share some high level strengths and weaknesses in terms of the key privacy design concepts we identified in Step Two.

Based on our analysis we found that the technologies that are being researched and developed can be used to meet three of the five key privacy design concepts identified. In considering concept 1, we compared the user scenarios that we collected from the organizations that participated in our interview research and the range of privacy solutions that we found in the literature. We did not find one solution that obviously met all of the users needs for providing a single solution that would protect data within the large organizations' highly heterogeneous and widely distributed configurations. Nor does it seem likely that one could be designed anytime soon. However, there are at least two approaches to addressing this problem. One approach is the creation of a common set of privacy utilities that allows users a single method for creating, visualizing and auditing privacy policies that could then be enforced using the appropriate range of technologies. Another possible approach is for a set of utilities to be provided to a central PI store on a single platform that has a privacy policy enforcement mechanism. This would create a PI "vault". Other distributed applications would then request data from that system. We recognize that there are privacy enabling technologies that address concepts 2 and 3. Many of the privacy approaches that have been identified allow the privacy enforcement to be separated from the application. For example, the Hippocratic Database [2] allows applications to query the database as they always have. The query re-writing done by the JDBC layer ensures that only PI accesses or updates allowed by the policy occur. Likewise, data store monitoring approaches such as that employed by Tivoli Privacy Manager [19] separate the application from the privacy auditing and/or enforcement. Each of these approaches also has the potential to allow privacy enforcement at the database field level.

Although we found approaches that can address the first three key privacy concepts, we have not found an approach that addresses either of the last two concepts. In the case of concept 4, the representatives of the organizations that we interviewed told us that they needed to be able to provide privacy protection for information within unstructured documents. Perhaps text analytics research combined with a privacy enforcement mechanism may be able to address this need in the future. Finally, none of the privacy technologies we analyzed addressed the last key privacy design concept (concept 5) that we identified. Organizational users have a need for simple and flexible interaction methods that can be used by individuals who do not have IT skills particularly in the areas of privacy policy creation and auditing. This need is consistent with the challenge of providing users with security and privacy controls they can understand that was identified by the CRA Conference on Grand Research Challenges in Information Security and Assurance [15]. This is also a need that can be addressed through the use of human computer interaction research and the application of user-centered design methods. Therefore, this is the need that we decided to address in our research. We identified three areas where highly usable privacy utilities were needed. The first is a utility to assist users in creating and understanding privacy policies. The second is a utility to assist users in implementing the privacy policy. The design of this utility is partially dependent on the choice of enforcement engines used. Finally the third utility enables organizations to conduct internal audits of their privacy policies. While our research has focused on all three areas, our work in the privacy policy creation area is the most mature and is the least dependent on a particular enforcement engine. Therefore, we will concentrate on this utility in this paper.

During the survey and interview research, many of the participants indicated that privacy policies in their organizations were created by committees made up of business process specialists, lawyers and security specialists as well as information technologists. Based on the range of skills generally possessed by people with these varied roles, we hypothesized that different methods of defining privacy policies would be necessary. Figure 2 shows the abstract architecture updated with a more detailed privacy policy creation utility. The figure shows the privacy policy creation utility divided into three parts. There is a privacy policy authoring utility that uses and stores natural language policies, a transformation utility for translating the policy into machine readable policies, and a visualization utility for helping users understand the implications of new and existing policies. The architectural view of this utility was used to guide the design of a prototype privacy management tool in Step Four.

Step Four: Designing and Evaluating a Privacy Policy Prototype

Using the completed survey and interview research and the architectural analysis discussed in Step Three, we designed and developed a prototype of a privacy policy management tool called SPARCLE (Server Privacy ARchitecture and CapabiLity Enablement). The overall goal in designing SPARCLE was to provide organizations with tools to help them create understandable privacy policies, link their written privacy policies with the implementation of the policy across their IT configurations, and then help them to monitor the enforcement of the policy through internal audits. Once we designed a prototype, we conducted a series of walkthrough sessions in which we utilized the prototype to discuss an appropriate scenario with representatives of health care, government, and finance organizations. In this paper, we will concentrate on the techniques we designed and developed for authoring privacy policies and assisting organizations in understanding the policies that have been created.

Authoring Privacy Policies

Based on the architectural drawings above, SPARCLE was designed to support users with a variety of skills by allowing individuals responsible for the creation of privacy policies to define the policies using natural language or to use a structured format to define the elements and rule relationships that will be directly used in the machine readable policy. SPARCLE keeps the two formats synchronized. For users who prefer authoring with natural language, SPARCLE transforms the policy into a structured form so that the author can review it and then transforms it into a machine readable format such as EPAL [7], XACML [29] or other appropriate privacy languages. SPARCLE translates the policies of organizational users who prefer to author rules using a structured format into both a natural language format and the machine readable version. During the entire privacy policy authoring phase, users can switch between the natural language and structured views of the policy for viewing and editing purposes. Once the machine readable policy is created, it is possible to employ any enforcement engine that is capable of using the elements of the standardized privacy policy language to ensure the policy is enforced for data stored in the organization's on-line data stores.

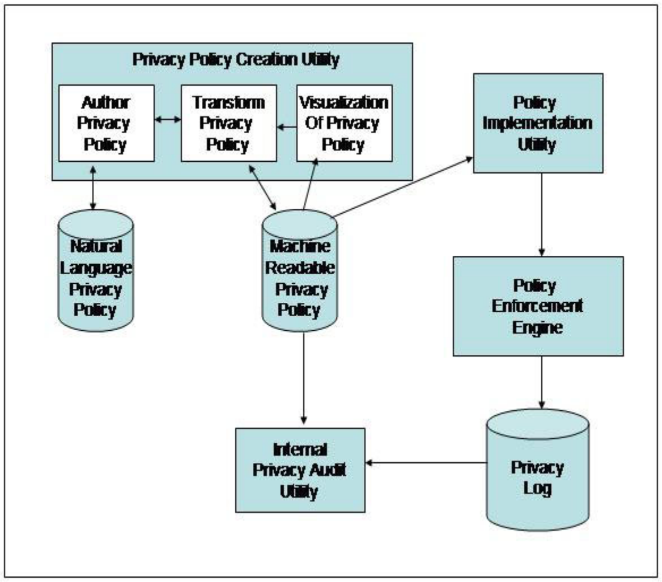
Privacy Policy Creation Utility
Author Privacy Policy
Transform Privacy Policy
Visualization Of Privacy Policy
Policy Implementation Utility
Natural Language Privacy Policy
Machine Readable Privacy Policy
Policy Enforcement Engine
Internal Privacy Audit Utility
Privacy Log

Figure 2 .Abstract Privacy Architecture with Privacy Creation Utilities Expanded

Figure 3 contains a screen capture of SPARCLE's natural language interface for defining privacy policies. Throughout SPARCLE, the tool provides a task flow in the form of tabs showing the high level task steps to be accomplished and the status of each. The tasks include: Author Policy (step shown in Figures 3), Transform Policy (step shown in Figure 4), Map User Categories, Map Data Categories, Map Purposes/Actions, Map Conditions, Map Obligations, and Verify Policy. The mapping steps are used to associate policy elements with system objects, and enable separation of high level and detailed policy specification and the Verify step allows users to confirm that all parts of the policy have been mapped. The page also contains general information about the policy, (the name, date created, and file source of the policy, and a description of the policy authoring task to be performed) a list of privacy policy templates that could be either provided by the tool for particular domains and geographies based on laws or created by the organization for customization and use by its divisions, and an Example Rule Guide describing the elements that make up a privacy policy rule. The guide is based on analyses of privacy policy rules specified in [7].

The guide defines the basic components that are necessary in a privacy policy rule that is enforceable including user categories, allowed actions, data categories, purposes, as well as optional components such as conditions and obligations. Finally, a text entry area is provided for the actual privacy policy. When the user begins the process of creating a new policy, she can create the policy from scratch by typing into the text entry area, copying an existing policy into that area, or selecting one of the templates provided and modifying the chosen rule template.

When the author is satisfied with the policy, he clicks on the Transform Policy tab shown in Figure 3. The natural language policy is analyzed and the policy elements (the strings which describe the User Category, Action, Data Category, Purpose, Conditions, and Obligations) in each rule are identified using a natural language parser. The natural language entry field area is replaced with a structured privacy policy creation view (shown in Figure 4).

The page also contains the policy information and the list of policy templates that was available on the policy authoring page. Next, the user is provided with a list containing the parsed rules in the current policy.

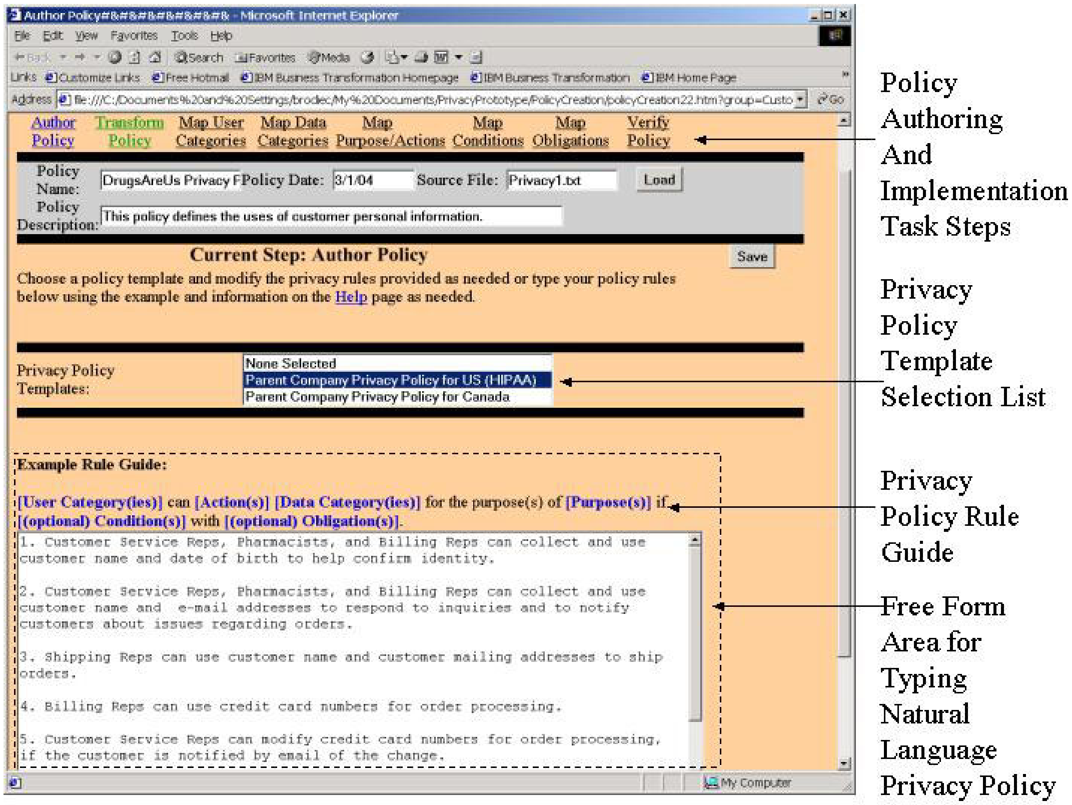
Author Policy#&#&#&#&#&#&#& - Microsoft Internet Explorer
File Edit View Favorites Tools Help
Back Search Favorites Media
Links Customize Links Free Hotmail IBM Business Transformation Homepage IBM Business Transformation IBM Home Page
Address file:///C:/Documents%20and%20Settings/brodiec/My%20Documents/PrivacyPrototype/PolicyCreation/policyCreation22.htm?group=Custo
Go
Author Policy
Transform Policy
Map User Categories
Map Data Categories
Map Purpose/Actions
Map Conditions
Map Obligations
Verify Policy
Load
Save
My Computer
Policy Authoring And Implementation Task Steps
Privacy Policy Template Selection List
Privacy Policy Rule Guide
Free Form Area for Typing Natural Language Privacy Policy

Policy Name: DrugsAreUs Privacy F
Policy Date: 3/1/04
Source File: Privacy1.txt

Policy Description: This policy defines the uses of customer personal information.

Current Step: Author Policy

Choose a policy template and modify the privacy rules provided as needed or type your policy rules below using the example and information on the Help page as needed.

Privacy Policy Templates:
None Selected
Parent Company Privacy Policy for US (HIPAA)
Parent Company Privacy Policy for Canada

Example Rule Guide:

[User Category(ies)] can **[Action(s)] [Data Category(ies)]** for the purpose(s) of **[Purpose(s)]** if **[(optional) Condition(s)]** with **[(optional) Obligation(s)]**.

```
1. Customer Service Reps, Pharmacists, and Billing Reps can collect and use
customer name and date of birth to help confirm identity.

2. Customer Service Reps, Pharmacists, and Billing Reps can collect and use
customer name and  e-mail addresses to respond to inquiries and to notify
customers about issues regarding orders.

3. Shipping Reps can use customer name and customer mailing addresses to ship
orders.

4. Billing Reps can use credit card numbers for order processing.

5. Customer Service Reps can modify credit card numbers for order processing,
if the customer is notified by email of the change.
```

Figure 3. SPARCLE Natural Language Privacy Policy Creation Screen.

Whenever a parsed rule is selected in the transformed view, the original unparsed text is also displayed and the elements of the rule that have been identified are highlighted in individual policy element selection lists as shown in Figure 4. There is one policy element selection list for each of the 6 types of rule element. There were two original purposes for this part of the prototype. First, while the natural language parsing technology in a limited domain such as privacy policy creation has promising accuracy, it is not perfect. This page allows users who have created the policy using the natural language technique to confirm that the parsing technology has identified all parts of the rules correctly and to correct anything that is in error. Second, for users who prefer the more structured method for privacy policy creation, this method can be used to create the entire policy. The organization or user can define policy element lists and then rules can be created by selecting the appropriate elements from each of the policy element selection lists and selecting “Add Rule”. Likewise, a rule can be modified or deleted by highlighting the rule in the rule selection list, modifying the selected elements as appropriate and selecting “Modify Rule” or “Delete Rule”. Any modification to rules or rules added or deleted using the structured approach is automatically reflected in the natural language version of the rules as well. Therefore, the author is able to go back and forth between the two methods to view the policy either in natural language or the parsed format with the elements identified.

During the course of the scenario-based sessions with users, they identified an additional use of the combined natural language and structured methods. The users indicated that it would be valuable to them for assessing the completeness of their existing privacy policies. Several participants were excited about the possibility of using SPARCLE to analyze their existing natural language privacy policies and then view the elements and rules identified in order to identify gaps and inconsistencies in the policies. For example, if an existing privacy policy rule fails to identify the purpose for which a particular user group is allowed to use a particular piece of data, the parsed rule would contain "none found" where purpose would usually be. The organizational users felt that this would be a valuable tool to ensure the quality of the privacy policies used by the organization and helpful in educating their organizations regarding their privacy policies.

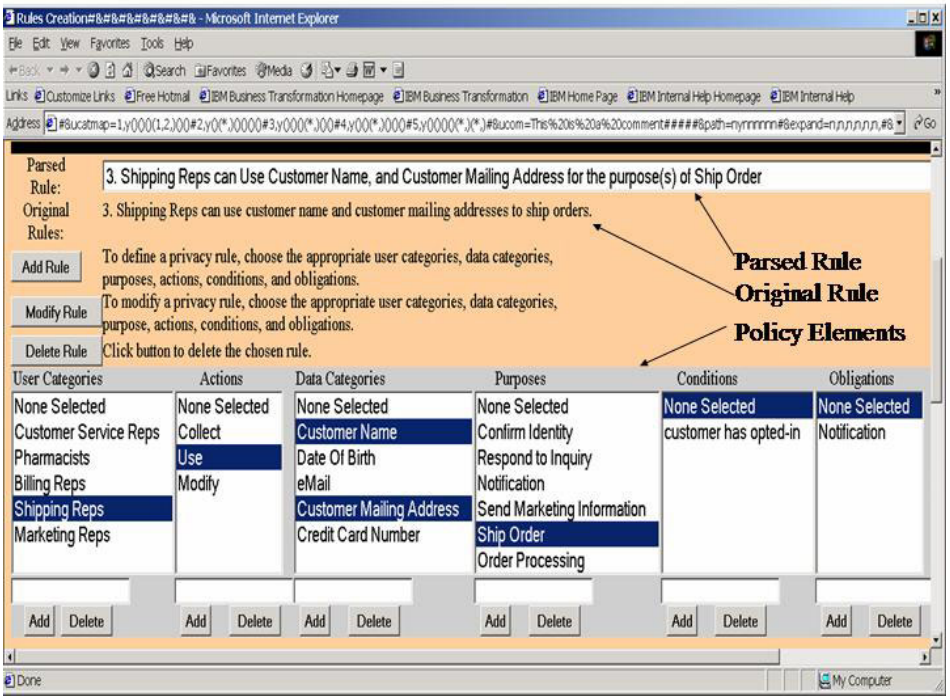
Rules Creation#&#&#&#&#&#&#& - Microsoft Internet Explorer
File Edit View Favorites Tools Help
Back Search Favorites Media
Links Customize Links Free Hotmail IBM Business Transformation Homepage IBM Business Transformation IBM Home Page IBM Internal Help Homepage IBM Internal Help
Address #&ucatmap=1,y000(1,2,)00#2,y0(*,)0000#3,y000(*,)00#4,y00(*,)000#5,y0000(*,)(*,)#&ucom=This%20is%20a%20comment#####&path=nynnnnn#&expand=n,n,n,n,n,n,#&
Go
Add Rule
Modify Rule
Delete Rule
Add
Delete
Done
My Computer

Parsed Rule: 3. Shipping Reps can Use Customer Name, and Customer Mailing Address for the purpose(s) of Ship Order

Original Rules: 3. Shipping Reps can use customer name and customer mailing addresses to ship orders.

To define a privacy rule, choose the appropriate user categories, data categories, purposes, actions, conditions, and obligations.

To modify a privacy rule, choose the appropriate user categories, data categories, purpose, actions, conditions, and obligations.

Click button to delete the chosen rule.

Parsed Rule

Original Rule

Policy Elements

User Categories	Actions	Data Categories	Purposes	Conditions	Obligations
None Selected	None Selected	None Selected	None Selected	None Selected	None Selected
Customer Service Reps	Collect	Customer Name	Confirm Identity	customer has opted-in	Notification
Pharmacists	Use	Date Of Birth	Respond to Inquiry		
Billing Reps	Modify	eMail	Notification		
Shipping Reps		Customer Mailing Address	Send Marketing Information		
Marketing Reps		Credit Card Number	Ship Order		
			Order Processing		

Figure 4. Expanded View of SPARCLE Structured Privacy Policy Rule Creation.

Understanding Privacy Policies

Based on the data collected from interviews with organizational users responsible for the creation of privacy policies, they often find it difficult to understand the policies that they create in order to ensure that policies are complete, able to be implemented, and consistent. Figure 5 shows our design to provide users with easy ways of viewing the privacy policy. The Figure contains a table in which two of the policy element types are used as axes and the other privacy rule elements that are associated with each row and column are shown in the cells. In the example that is shown, user categories are placed on the horizontal axis and data categories are placed on the vertical axis. The cells in the table contain the purposes, conditions, and obligations for rules that apply to that user and data category. Using this table, users can see at a glance what type of users are allowed to access each data element and also see which user groups are never allowed to access particular data items. While the table format was well received by users, we are not yet sure how well a two dimensional table scales up to real organizational policy complexity. Scaling and visualization will be the subject of our future research.

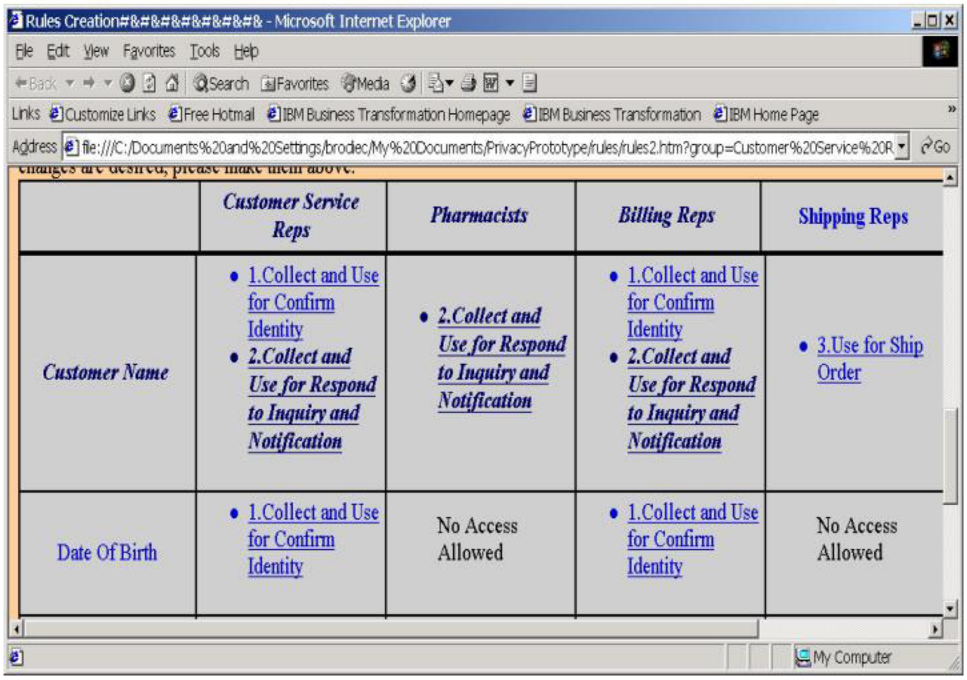
Rules Creation#&#&#&#&#&#& - Microsoft Internet Explorer
File Edit View Favorites Tools Help
Back Search Favorites Media
Links Customize Links Free Hotmail IBM Business Transformation Homepage IBM Business Transformation IBM Home Page
Address file:///C:/Documents%20and%20Settings/brodiec/My%20Documents/PrivacyPrototype/rules/rules2.htm?group=Customer%20Service%20R
Go
My Computer

changes are desired, please make them above.

	Customer Service Reps	*Pharmacists*	*Billing Reps*	Shipping Reps
Customer Name	• 1.Collect and Use for Confirm Identity • *2.Collect and Use for Respond to Inquiry and Notification*	• *2.Collect and Use for Respond to Inquiry and Notification*	• 1.Collect and Use for Confirm Identity • *2.Collect and Use for Respond to Inquiry and Notification*	• 3.Use for Ship Order
Date Of Birth	• 1.Collect and Use for Confirm Identity	No Access Allowed	• 1.Collect and Use for Confirm Identity	No Access Allowed

Figure 5. Table Showing Privacy Policy Rules that Apply to Each User and Data Category.

Results from Design Feedback Sessions with Privacy Professionals

We conducted scenario-based usability walkthrough sessions of two iterations of SPARCLE with people who were responsible for the creation, implementation, and auditing of privacy policies within large organizations in the domains of healthcare, banking, and government. During the course of the 90 minute sessions with 1 to 4 participants, we gathered verbal and written feedback on the usability, design, and value of the privacy tool. The summary results are presented in Figure 6. This figure shows the results for the features that were most highly rated by the walkthrough participants. Participants in iteration 1 requested the ability to work with a set of pre-loaded templates for domains or organizations, and the participants in iteration 2 rated this very highly. In the second iteration the ability to enter rules with natural language with a guide and the structured method for rule creation were rated essentially the same by participants. Other highly rated features included the table visualization of the policy which provided users with a big picture view of the coverage of a policy, the ability to create the policy rules first and then later map them to elements in data stores, and the ability to run queries on the purpose based accesses to particular individuals' personal information.

For the first iteration of the prototype, walkthrough participants (7 participants in 5 sessions) rated the prototype positively (an average rating of 5.39 on a 7-point scale with 1 indicating "no value" and 7 indicating "highest value"). We present this summary result since it communicates the overall positive response to the prototype. However, the primary purpose for the sessions was to gather more qualitative responses from the participants about the value of the system to their task of managing privacy.

At the conclusion of the first iteration of design and evaluation, we made the following changes: 1) We added the ability to import pre-existing privacy policies into the natural language policy authoring condition to allow SPARCLE to highlight gaps and inconsistencies in the policies, 2) We added the ability to use privacy policy templates as a starting point for authoring privacy policies using either the natural language or structured policy authoring methods, and 3) We improved the readability of the table view of the privacy policy by bulletizing entries and making it scrollable. During the second iteration of walkthrough sessions, the participants (15 participants in 6 sessions) also rated the revised prototype very positively (an average rating of 5.55 on the same scale).

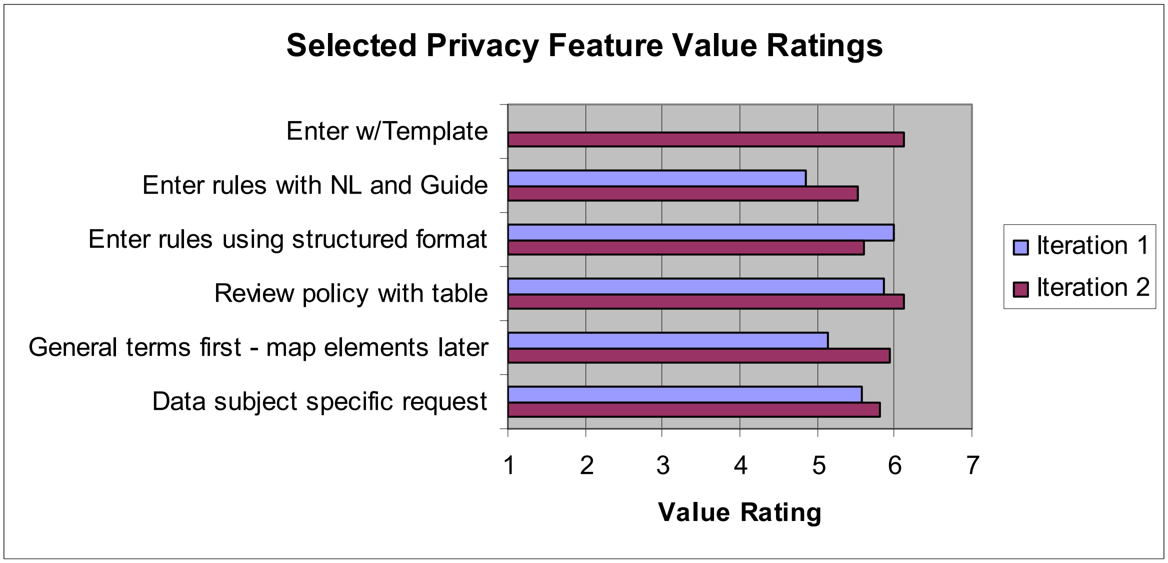

Figure 6. Privacy Management Features with Highest Participant Ratings from Design Walkthrough Sessions

Step Five: Evaluating Policy Authoring

An empirical laboratory study was run to compare the two privacy policy authoring methods illustrated in the prototype. In order to provide a baseline comparison for the two methods (Natural Language with a Template, and Structured Entry from Element Lists), we added a control condition that allowed users to enter privacy policies in any format that they were satisfied with (Unconstrained Natural Language).

Experimental Design

Thirty-six employees of a large IT company were recruited through email to participate in the study. The participants had no previous experience in privacy policy authoring or implementation. A repeated measures design was employed in the study; each participant completed one task in each of the three conditions. All participants started with a privacy rule task in the Unconstrained control condition (Unconstrained). Then, half of the participants completed a similar task in the Natural Language with a Template condition (Template), followed by a third task in the Structured Entry from Element Lists condition (Structured). The other half of the participants completed the Structured condition followed by the Template condition.

In each task, we instructed participants to compose a number of privacy rules for a pre-defined task scenario. Participants worked on three different scenarios in the three tasks. We developed the scenarios in the context of three privacy sensitive domains, namely health care, government, and banking. Each scenario contained five or six privacy rules, including one condition and one obligation. The order of the scenarios was balanced across all participants. We recorded the time that the participants took to complete each task and the privacy rules that participants composed. We also collected, through questionnaires, participants' perceived satisfaction with task completion time, quality of rules created, and overall experience after participants completed each task. At the end of the session, participants completed a debrief questionnaire about their experiences with the three rule authoring methods.

In order to compare the quality of the rules participants created under different conditions and scenarios, we developed a standard metric for scoring the rules. We counted each element of a rule as one point. Therefore, a basic rule of four compulsory elements had a score of four and a scenario that consisted of five rules, including one condition and one obligation, had a total score of 22. We counted the number of correct elements that participants specified in their rules, and divided that number by the total score of the specific scenario. This provides a standardized score of the percentage of elements correctly identified that was compared across different scenarios and conditions.

Results and Discussion

There was a significant difference in the task completion time across the three conditions ($F_{(2, 70)} = 4.58$, $p < 0.05$). Mean participant time on task was 910 seconds for Unconstrained, 814 for Template, and 992 for Structured conditions respectively. Post hoc tests showed that the Template method took significantly shorter time than the Structured method. There was no significant difference between the Unconstrained method and the other two methods.

A repeated measures test with post hoc analyses indicated that participants were more satisfied with the quality of the rules created by the Template method or the Structured method as compared with the Unconstrained method ($F_{(2, 70)}) = 6.54$, $p < 0.005$). On a scale of 1 to 7, with 1 indicating highest overall satisfaction, participants mean satisfaction scores were 4.0 for Unconstrained, 3.3 for Template and 3.4 for Structured conditions. There was no significant difference between the Template method and the Structured method.

A statistical test of the rule quality scores calculated using the standard metric found a significant difference between the three conditions ($F_{(2, 70)} = 44.3$ $p < 0.001$) (see Figure 7 below). Post hoc tests showed that the Template method and the Structured method helped users create rules with significantly higher quality than the Unconstrained method. There was no significant difference between the Template method and the Structured method. Using the Unconstrained method, participants correctly identified about 42% of the elements in the scenarios, while the Template method and the Structured method helped users to correctly identify 75% and 80% of the elements, respectively. Considering the fact that the participants were novice users, some of the improvement might have been the result of learning in the first trial. However we did not provide feedback on rule quality so attribute most of the differences in performance to the authoring methods.

We examined the data in more depth by separating the total number of elements identified by participants into six privacy policy element categories, to examine whether there were any categories that caused particular difficulties (see Figure 8). A two way repeated measures test with post hoc analysis found that the Template method and the Structured method significantly outperformed the Unconstrained method across all categories except the condition category. There was no significant difference between the Template condition and the Structured condition for any of the 6 categories. For the Unconstrained condition, we found that participants were more likely to identify the data element than the data user ($t_{(35)} = -4.45$, $p < 0.001$), suggesting that novice users tended to focus on 'what data can be accessed' rather than 'who should have access to the data'.

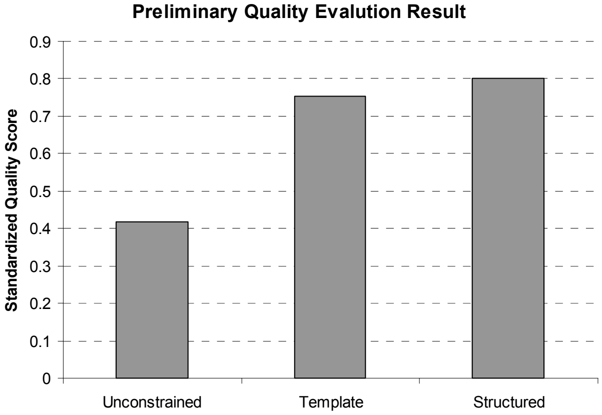

Figure 7. Average Scores of the Quality of the Rules According to the Quality Evaluation Metric in Three Conditions.

The results of the experiment confirmed that both the Template method and the Structured method enabled participants to create rules with higher quality than the Unconstrained method. And the fact that the percentages of elements identified by the Template method and the Structured method almost doubled that of the Unconstrained method suggests that the Template and the Structured methods are reasonably easy to learn and use.

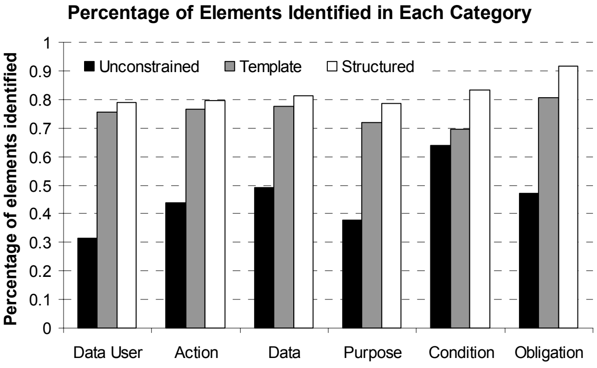

Figure 8. Average Percentages of the Elements Identified in Each Data Category in Three Conditions.

CONCLUSIONS AND FUTURE RESEARCH

Privacy is emerging as a powerful issue for people within organizations and individuals who interact with them around the world. In the networked world in which we live today, the topic is of growing concern and importance and highly usable technology is needed to enable protection of personal information. Previous research has shown that the general public is concerned about protecting their privacy [20] and that many organizations do not feel they are adequately protecting PI [32]. Our research contributes both to the development of an understanding of privacy and to identifying methods for developing usable security and privacy technologies. This research provides insight about the developing privacy perspectives, concerns, and needs of organizations. The case study highlights the work of identifying organizational privacy requirements, iteratively designing and validating a prototype with users in their work settings, and empirical laboratory testing to guide specific design decisions to meet the needs of providing flexible privacy technologies for organizations and their users. The research illustrates the selection and use of a variety of user-centered methods. The methods that we selected at each point in the research were tailored to the context at different points in the evolution of the research topic.

We began by selecting and developing an email survey instrument to assist in understanding the current state of privacy in organizations. The email survey data in Step One highlighted potential emerging trends in the views of organizational representatives in North America, Europe, and Asia towards privacy. We found that there were differences based on organizational domain and geography, but also found considerable agreement on the need to better integrate policy and tools. Next we used in-depth interviews to assist us in understanding typical flows of personal information within organizations. In Step Two we developed scenarios for use in ongoing discussions with organizations concerning PI data flows. By utilizing a mixture of survey and interview research in the first two steps, we felt that we efficiently developed an understanding of privacy management to use as a focus in the design of privacy management tools. We learned that organizations are concerned about protecting PI both at rest and in transit from both misuse by internal users as well as external sources, linking privacy policies to their IT configurations, and about ensuring that their policies are being enforced correctly.

Building on the results of this work, we conducted an architectural analysis. Using the requirements described above, we analyzed a range of existing and proposed privacy solutions to understand how well they met the privacy needs of organizations in Step Three. We found that although there are a wide variety of technologies that could be used to protect the PI held by organizations, they are not currently being widely used because they are not designed with the skills of potential users in mind. Based on this we identified a set of privacy utilities that are needed to allow organizations to create understandable policies, link their written privacy policies with the implementation of the policy across their IT configurations, and then help them to monitor the enforcement of the policy through internal audits. These utilities were prototyped in the user-centered design and evaluation method we used in Step Four. We explored and iterated on the design with users and were able to obtain valuable feedback well before we could complete a full implementation of the prototype. While work on the natural language parsing and mapping components of SPARCLE is still underway, we think we have gained a solid understanding of organizational requirements for its successful completion.

We augmented our design work by carrying out a controlled study of the authoring methods we developed for the prototype. In Step Five we conducted an empirical usability laboratory test of the structured list entry and natural language methods for authoring policies. Results were promising and showed that in initial use, novice users could use the two methods to identify and cover 75-80% of the policy rule elements. Coupled with our work with organizational users, we have concluded that integrating the Structured and Template authoring methods along with providing an easy to understand policy coverage view will be important elements of a successful privacy policy tool. We think that the laboratory test was an important component of the overall research in helping to justify the value and improve the overall design of each authoring approach.

We think that a number of research challenges remain. First, we need to examine how well our authoring environment works for realistically complex organizational privacy policies. Our users have generally been from large organizations, and they have responded well to the parts of the prototype we present in this paper – authoring and viewing policy coverage. However, working with policies with hundreds of rules might create problems that do not emerge in discussions centered on a single policy involving a few rules. A planned next step and a natural evolution for our research will be to work with several organizations to create complete machine readable policies which reflect their actual internal privacy policies.

While the results of our research into understanding and addressing organizational user needs for privacy will be useful to organizations in helping them protect the privacy of their customer, constituents, patients, and employees, we believe that a secondary value of this work is as an example of how to create more usable security software. Multiple researchers [15, 39] have identified usability as one of the grand challenges for security and privacy research. The application of user-centered methods and HCI research techniques described in this paper could serve as a model for the design of interaction methods for many security projects. At the same time, current world events are providing pressure on the public and private sectors to better protect both the privacy and security of the data held by all types of organizations. HCI research and the application of user-centered design techniques can help the security and privacy community step up to the challenge of creating interfaces and interaction methods that reduce the complexity in defining, implementing, and managing privacy policies and security solutions for the benefit of all parties.

REFERENCES

1. Adams, A. and Sasse, A. (2001) Privacy in Multimedia Communications: Protecting Users, Not Just Data. In A. Blandford, J. Vanderdonkt & P. Gray [Eds.]: People and Computers XV - Interaction without frontiers. *Joint Proceedings of HCI2001 and ICM2001*, Lille, Sept. 2001. pp. 49-64. Springer.
2. Agrawal, R., Kiernan, J., Srikant, R., and Xu, Y. Hippocratic Databases. *Proceedings of the 28th Very Large Database Conference (VLDB),* Hong Kong, China, 2002.
3. Agrawal, R., Kiernan, J., Srikant, R., and Xu, Y. Implementing P3P Using Database Technology. *Proceedings of the 19th International Conference on Data Engineering*, Bangalore, India, 2003.
4. Anderson R. J. A Security Policy Model for Clinical Information Systems. *In the Proceedings of the 1996 IEEE Symposium on Security and Privacy*, 30-43.
5. Anderson R. J. Privacy Technology Lessons from Healthcare. *In the Proceedings of the 2000 IEEE Symposium on Security and Privacy.*
6. Anton, A., He, Q., and Baumer, D. (2004) The complexity underlying JetBlue's privacy policy violations. *IEEE Security & Privacy*. August/September, 2004.

7. Ashley, P., Hada, S., Karjoth, G., Powers, C., and Schunter, M. (2003). *Enterprise Privacy Architecture Language (EPAL 1.2).* W3C Member Submission 10-Nov-2003. http://www.w3.org/Submission/EPAL/

8. AT&T Privacy Bird (2003). http://privacybird.com/

9. Backes, M., Pfitzmann, B., and Schunter, M. A Toolkit for Managing Enterprise Privacy Policies. *In the Proceedings of the 8th European Symposium on Research in Computer Security (ESORICS)*, Springer-Verlag, Berlin, 2003.

10. Ball, E. (2003). Patient privacy in electronic prescription transfer. *IEEE Security and Privacy*, 1, 2, 77-80.

11. Baumer, D., Earp, J.B., and Payton, F. C. (2000). Privacy in medical records: IT implications of HIPAA. *Computers and Society*, December, 2000, 40-47.

12. Beyer, H. and Holtzblatt, K. (1988). *Contextual Design*. NY: Morgan Kaufmann.

13. Bohrer, K., Levy, S., Liu, X., and Schonberg, E. Individual Privacy Policy Based Access Control. In Proceedings of the 6th International Conference on Electronic Commerce Research (ICECR-6), October, 2003, Dallas, Texas.

14. Chatham, B. (2004). Online Privacy Concerns: More than Hype. *The Forrester Report*, 2004

15. CRA Conference on "Grand Research Challenges in Information Security and Assurance". http://www.cra.org/Activities/grand.challenges/security/. November 16-19, 2003.

16. Cranor, L. (2002). *Web Privacy with P3P.* Cambridge: O'Reilly.

17. Hagen, P. (2000). Personalization versus privacy. *The Forrester Report*, Nov., 2000, 1-19.

18. IBM. *IBM Multi-National Consumer Privacy Survey.* http://www.ibm.com/services/files/privacy_survey_oct991.pdf

19. IBM Tivoli Privacy Manager for eBusiness (2004). http://www-306.ibm.com/software/tivoli/products/privacy-mgr-e-bus/.

20. Jensen, C. and Potts, C. (2004). Privacy polices as decision-making tools: An evaluation of online privacy notices. *CHI 2004*, 471-478.

21. Karat, C. Iterative Usability Testing of a Security Application. *In Proceedings of the Human Factors Society 33rd Annual Meeting*, 1989.

22. Karat, C., Brodie, C., and Karat, J. (2003). Views of Privacy: Business Drivers, Strategies, and Directions. *IBM Research Report RC22912 (W0309-132).*

23. Karat, C., Brodie, C., Karat, J., Vergo, J., and Alpert, S. (2003) Personalizing the user experience on ibm.com. *IBM Systems Journal*, 42, 4, 686-701.

24. Karjoth, G. and Schunter, M.(2002) A Privacy Policy Model for Enterprises. *Proceedings of the 15th IEEE Computer Security Foundations Workshop*, 271-281.

25. Kobsa, A. (2002) Personalized hypermedia and international privacy. *Communications of the ACM*, 45, 5, 64-67.

26. Manny, C. H. (2003). European and American privacy: Commerce, rights, and justice. *Computer Law and Security Report*, 19, 1, 2003, 4-10.

27. Microsoft Internet Explorer (2004). Help Safeguard your privacy on the web. http://www.microsoft.com/windows/ie/using/howto/privacy/config.mspx.

28. National Research Council. (2003). *Who goes there? Authentication through the lens of privacy*. Washington, D.C: National Academies Press.

29. OASIS (2004). Privacy Policy Profile of XACML draft 01. http://docs.oasis-open.org/xacml/access_control-xacml-2_0-privacy_profile-spec-cd-01.pdf

30.OECD (1980). OECD guidelines on the protection of privacy and transborder flows of personal data. http://www.oecd.org/home/

31.Office of the Federal Privacy Commissioner of Australia. (2000). *Privacy and Business (2000).* http://www.privacy.gov.au.

32. Ponemon Institute and IAPP. (2004). 2003 Benchmark Study of Corporate Privacy Practices.

33.Senior, A., Pankanti, S., Hampapur, A., Brown, L., Tian, Y., and Ekin, A. (2004). Blinkering Surveillance: Enabling Video Privacy through Computer Vision. *IEEE Security and Privacy*, in press.

34.Shneiderman, B. and Plaisant, C. (2004). *Designing the User Interface*. Reading: Addison Wesley.

35.Smith, J. (1993). Privacy policies and practices: Inside the organizational maze. *Communications of the ACM*, 36, 12, 105-122.

36.The Stanford Student Computer and Network Privacy Project. A study of student privacy issues at Stanford University. *Communications of the ACM*, 45, 3, 2002, 23-25.

37.U.S. Fair and Accurate Credit Transaction Act. (2003). H.R. 2622, 108th Congress, July 24, 2003.

38.Warren, S.A. and Brandeis, L.D. (1890). The right to privacy. *Harvard Business Review*, Dec, 4, 195

39.Whitten, A. and Tygar J. D. (1999) Why Johnny Can't Encrypt: A Usability Evaluation of PGP 5.0. *In Proceedings of the 9th USENIX Security Symposium*, August, 1999.

40.W3C (2002) A P3P Preference Exchange Language 1.0 (APPEL 1.0). http://www.w3.org/TR/P3P-preferences/

41. Zurko, M. E., Simon, R., and Sanfilippo, T. (1999) A User-Centered, Modular Authorization Service Built on an RBAC Foundation. *In Proceedings of the 1999 Symposium on Security and Privacy*, May, 1999, 57-71.