Specifying Privacy Policies with P3P and EPAL: Lessons Learned

William Stufflebeam, Annie I. Antón, Qingfeng He, and Neha Jain
Department of Computer Science, North Carolina State University, Raleigh, NC 27695
{whstuffl, aianton, qhe2, njain}@eos.ncsu.edu

Abstract

As computing becomes more ubiquitous and Internet use continues to rise, it is increasingly important for organizations to construct accurate and effective privacy policies that document their information handling and usage practices. Most privacy policies are derived and specified in a somewhat ad-hoc manner, leading to policies that are of limited use to the consumers they are intended to serve. To make privacy policies more readable and enforceable, two privacy policy specification languages have emerged, P3P and EPAL. This paper discusses the effectiveness of these languages within the context of a case study that entailed the expression of common online privacy statements for a healthcare website, employing requirements engineering quality factors as a framework for our discussion.

Specifying Privacy Policies with P3P and EPAL: Lessons Learned

William Stufflebeam, Annie I. Antón, Qingfeng He, and Neha Jain
Department of Computer Science, North Carolina State University, Raleigh, NC 27695
{whstuffl, aianton, qhe2, njain}@eos.ncsu.edu

Abstract

As computing becomes more ubiquitous and Internet use continues to rise, it is increasingly important for organizations to construct accurate and effective privacy policies that document their information handling and usage practices. Most privacy policies are derived and specified in a somewhat ad-hoc manner, leading to policies that are of limited use to the consumers they are intended to serve. To make privacy policies more readable and enforceable, two privacy policy specification languages have emerged, P3P and EPAL. This paper discusses the effectiveness of these languages within the context of a case study that entailed the expression of common online privacy statements for a healthcare website, employing requirements engineering quality factors as a framework for our discussion.

Keywords: Policy Specification, P3P, EPAL, Healthcare, Privacy, Requirements Engineering

1. Introduction

Privacy policy specification and enforcement has become a hotbed of research activity over the past few years as Internet use has been on the rise around the globe. As the number of consumers participating in online activities grows, it becomes increasingly imperative for organizations to express their privacy practices in an accurate, accessible, and useful way. This enables consumers to evaluate an organization's information handling and usage practices before determining whether or not they wish to engage in transactions with the organization.

The Federal Trade Commission (FTC) defines a privacy policy as a comprehensive description of a company's information practices, accessible by clicking on a hyperlink on the company's website [FTC98]. Online privacy policies are thus documents made available by organizations via their websites that explain how consumers' personal information will be collected, used, and stored by the organization. These policies, which are the focus of this study, are intended to help consumers make informed decisions with regard to the organizations with which they interact and share their personal information.

Privacy policies are often ambiguous and difficult for the average Internet user to read, making them less useful to the website's visitors than is desirable [AEB04]. One reason for these difficulties is the lack of standardization (e.g. with vocabulary), even in the face of legislation, such as the Gramm-Leach-Bliley Act (GLBA)[1], which states that policies must be "clear and conspicuous" [AEB04]. Policies are not always accessible from a company's main homepage; this inconsistent location of privacy policies means that end-users must sometimes "hunt" for privacy practice information [BHA04]. Additionally, no matter the specific webpage with which a consumer may be interacting, he/she is always presented with the same policy information. The inability to obtain context-dependent privacy practice information places significant cognitive burden on the end-user. The lack of a standardized privacy policy vocabulary, the inconsistent location of privacy policies across websites, and the lack of context-dependent privacy practice information make the current state of privacy policy specification unsuitable for consumers [BHA04]. Two privacy policy standardization attempts have sought to address some of these problems, as we discuss herein.

The Platform for Privacy Preferences Project (P3P) is an attempt to provide a standardized, XML-based policy specification language that can be used to specify an organization's privacy practices in a way that can be parsed and used by policy-checking agents on the user's behalf [CDH04, Cra03, Hoc02, Jam04, KSH03b, Pre02].

P3P policy documents are semi-structured so that they can be parsed and evaluated against predefined preferences that a user indicates about how and when their information can be collected and handled. There are many user software agents available for use (e.g. Privacy Bird[2], Privacy Companion[3], Internet

[1] Gramm-Leach-Bliley Act of 1999, 15 U.S.C. §§ 6801-6809(2000).

[2] http://www.privacybird.com

[3] http:// www.idcide.com/

Explorer 6.0[4]), which handle policy checking and invoke the required actions that need to occur when a website's policy is found to conflict with the user's preferences. These actions range from blocking a particular webpage from being displayed, to placing a warning icon in the user's browser status bar. With this technology, P3P-conversant users can partially automate limited privacy policy evaluation in an attempt to better control how their information is being collected, stored, and used. P3P is typically characterized as supporting machine-readable policies; it does not provide an enforcement mechanism for organizations to use in monitoring their information handling practices.

IBM's Enterprise Privacy Authorization Language (EPAL) addresses the need for machine-enforceable policies [AHK03]. Like P3P, EPAL is an XML-based privacy policy specification language, but it is designed for organizations to specify internal privacy policies. These EPAL policies can be used internally and amongst the organization and its business partners to ensure compliance with the underlying policies of each [PAS02, Sch03, SHW04].

To develop a better understanding of P3P's and EPAL's strengths and limitations, we conducted a case study in which we employed both specification languages to specify the privacy policies for the CIGNA HealthCare website [Cig04a, Cig04b]. The intent of this study is not to compare P3P with EPAL, as the two languages have different purposes and scopes, but to evaluate each language's expressiveness for specifying natural language privacy policies. The rest of this paper is organized as follows: Section 2 provides an overview relevant work. Section 3 describes our case study. Section 4 discusses lessons learned. Finally, Section 5 presents a summary and discusses future research directions.

2. Relevant Work

As privacy becomes an increasingly important concern to the general public, more research is examining privacy policy specification. This section summarizes the most relevant work to establish the necessary context for discussing our policy specification case study.

2.1 The Lack of Clarity in Privacy Policies

Studies show that most website privacy policies lack the clarity necessary for most Internet users to actually find them useful [AEB04, Hoc01]. Common policy clarity problems include the monolithic, context-independent nature of most policies [BHA04], the legalese commonly used to write the policies [AEB04], and the redundancies, ambiguities, and conflicts found in many policies [AER01, AEB04].

Most policies are monolithic in nature (i.e. There is one large privacy policy document, or several large privacy documents, rather than many smaller, topic-specific policies), and they are generally context-independent. This means that no matter which page on a website with which a consumer may be interacting, he/she is always presented with the same policy information, a situation which places significant cognitive burden on the end-user. Bolchini *et al*. have developed a contextualization process for privacy policies that structures them as usable, focused "micro-policies" linked according to the relevant "interaction contexts" on the site [BHA04]. This contextualization allows users to view privacy policy information relevant to specific interactions and transactions. The contextualization approach was validated for three customer interaction scenarios on the Amazon.com website [BHA04]. This case study demonstrates the methodical application of a five-step process, which yields a more meaningful and useful contextualized website privacy policy.

Our previous analysis of financial privacy policies [AEB04] revealed that it is challenging (even for experienced privacy policy analysts) to understand what different policy statements mean. The study entailed goal-based analysis of 40 Internet privacy policies from nine GLBA covered financial institutions. For the study, we used a software engineering technique (goal mining) to extract goal statements from privacy policies in an effort to determine the policies' true intents. This study revealed that certain policies used their own distinct vocabularies, despite the existence of laws that advocate standardization (such as GLBA). This vocabulary difference requires end-users to spend a great deal of time recalibrating their understanding when evaluating different policies. Our research shows there is a need for institutions to standardize the way in which they express their privacy practices, and that it is possible to make online privacy policies more clear, benefiting both the institutions and end-users.

2.2 Policy Analysis and Specification in Requirements Engineering (RE)

Requirements engineering (RE) is the process of discovering real-world goals for which a software system is intended by identifying stakeholders and their needs, the functions of and constraints on the intended system, and documenting these in a form that is amenable to analysis, communication and

[4] http://www.microsoft.com/windows/ie/default.mspx

subsequent implementation [NE00]. Requirements and policies share some similarities: both express desire and worth, rather than fact, and both are optative statements, specifying what must be done [AEP01, AER01]. In the RE community, researchers are using RE techniques to analyze and specify security and privacy policies and related requirements. The work presented in this paper is yet another piece of work predicated on the application of RE principles in privacy policy specification.

Pfleeger proposed a framework for security requirements in which requirements are distinguished from the security policies that constrain them and the controls and mechanisms that implement them [Pfl91]. Moffett has explored the relationship between requirements and high-level security policies [Mof99] and has addressed the challenges inherent in integrating different requirements specification techniques in the analysis of safety and security requirements [EM99]. Liu *et al.* applied the *i** framework, a goal-based requirements analysis method [Yu93], to provide access control analysis by modeling the dependencies among actors, tasks, and system resources [LYM03]. Fontaine [Fon01] employs KAOS, a requirements acquisition and refinement approach [DLF93], to refine security requirements into specific authorization rules and access control policies expressed in Ponder. Ponder is a language for specifying management and security policies for distributed systems [Dam02]. Crook et al. proposed a requirements-level analytical role-modeling framework to model access control policies based on organizational theories [CIN03].

Antón et al. specifically focus on privacy policies and requirements. Their major contributions to date include (1) providing techniques to aid in aligning software requirements with security and privacy policies [AEC03], (2) employing GBRAM, a goal-based requirements analysis method, to analyze a set of online privacy policy statements across several domains (e.g., healthcare, financial, e-commerce) [AEB04, AHB04, BHA04], and (3) proposing a privacy requirements taxonomy to classify privacy statements as either protection goals or vulnerabilities [AER01]. Privacy protection goals express ways in which a system protects the privacy of a particular piece of information, while privacy vulnerabilities express potential privacy vulnerabilities that exist within a system.

2.3 Privacy Policy Specification

We now discuss P3P and EPAL; both are attempts to standardize privacy policy specification.

2.3.1. Overview of P3P

P3P is a semi-structured privacy policy specification language that allows an organization to specify its website privacy practices in a machine-readable format. These P3P specifications can be checked by a Web browser or user agent, against user-specified preferences, to determine whether the organization follows user-acceptable privacy practices. This process is sometimes automated through software in a question-answer format. As the user browses P3P-enabled websites, the user's agent checks pages that the user attempts to load against these preset preferences and takes appropriate actions, such as allowing the load, preventing the load, or notifying the user that the site does not (or may not) comply with the user's preferences. A P3P policy expresses the privacy practices related to the particular page or pages it governs; it covers any information collection on those pages, the purposes of that collection, the information recipient, and the length of that information's retention.

P3P also provides a mechanism for specifying cookie-related privacy practices via *compact policies* [CDH04]. These compact policies are optional and simply offer possibility for optimization by allowing user agents to check the privacy policies regarding a page's cookies without loading a separate policy document. Compact Policies are included in the HTTP response headers for a given webpage, and they allow a user agent to quickly assess the cookie-related policies of the page being loaded and respond accordingly. If an agent cannot glean the necessary information from the compact policy, then it may refer to the full P3P policy to find what it is looking for.

In addition to the individual P3P statements that describe the collection, use, recipient, and retention of specific data, P3P policies also contain meta-information. This information includes the P3P policy name, contact information for the organization (or a third party that handles disputes regarding the website for the organization), and an optional expiration date for the policy, which guarantees the user how long the policy can be assumed to be valid [CDH04]. A P3P statement block, which covers the collection, use, recipient, and retention of a specific type of data, contains the elements defined in Table 1.

Table 1: Elements of a P3P Statement [CDH04]

P3P Element	Description
Consequence	Optional in P3P; helps explain why the suggested practice may be valuable in a specific instance even if the user would not normally allow the practice.

Purpose	Contains one or more purposes describing how collected data will be used. Possible values can be current, admin, develop, tailoring, pseudo-analysis, pseudo-decision, individual-analysis, individual-decision, contact, historical, and telemarketing.
Recipient	Contains one or more entities who will receive the collected data. Possible values can be ours, delivery, same, other-recipient, unrelated, and public.
Retention	Indicates kind of retention policy that applies to data referenced. Possible values can be no-retention, stated-purpose, legal-requirement, business-practices, and indefinite.
Data-group	Describes the data that will be collected and transferred.

2.3.2. Overview of EPAL

EPAL is an IBM project that is being developed mainly as a business-to-business (B2B) technology to help streamline information flows during business interactions [AHK03]. It helps ensure that information is protected and used in accordance with the responsible organization's privacy policies. IBM introduced EPAL as a formal language that provides enterprises with a way to automate and enforce privacy policies across IT applications and systems. The language allows organizations to specify their privacy practices in a way that they, and other organizations with which they interact, can read and use. EPAL policies, unlike P3P policies, are enforceable, as they are written and structured in a similar fashion to access control policies that one may find in the security domain. The policies are enforced by an enforcement engine that parses the files, assuring the information collection, use and storage that ocurrs within the organization, and amongst the organization and its partners, complies with the EPAL specified privacy practices.

Like P3P, EPAL policies also contain meta-information that does not specifically address information access and usage. The meta-information includes the policy ID and description, information about the issuing organization, and modification dates/document revision numbers [AHK03]. In EPAL, organizations define a vocabulary specific to their needs, and the only resulting requirement is that every agent that wants to use the policy to govern their interactions must agree upon and understand the vocabulary being used. EPAL rules specify the policies regarding a specific information access. The elements of an EPAL rule are defined in Table 2.

2.3.3 Existing Research on P3P and EPAL

Byers *et al*. developed a tool, which they used to analyze the data practices of 588 P3P-enabled websites [BCK03]. This was the first comprehensive automated analysis of a wide spectrum of P3P–enabled websites' data practices and demonstrated the feasibility of automated statistical analyses. Additionally, Hochheiser's general overview of P3P, examines the privacy model underlying P3P within the political context in the United States, and critiques P3P from a technical and social perspective [Hoc02]. He identifies several potential technical problems with P3P that are relevant to this work, including the implications of its limited scope, the limits of the language's vocabulary, and the dependence of its success on the implementation of usable user agents.

Table 2: Key Elements of an EPAL Rule [AHK03]

Rule Element	Description
Ruling	Describes one authorization answer such as 'allow' or 'deny'.
User-Category	Defines the user (i.e. enterprise) categories that can access the data e.g., a primary care physician. User Categories are hierarchies. Can be zero or more for each vocabulary.
Action	Models how the collected data is used e.g., disclose vs. read. Lists which actions are allowed or denied for certain purposes and under defined conditional circumstances. Can be zero or more for each vocabulary.
Data-Category	Defines different categories of collected data used to distinguish different data classifications e.g., medical record vs. contact data. Data categories are hierarchies. Can be zero or more for each vocabulary.
Purpose	Represents the intention for which information is disclosed and serves as a rule that authorizes or denies access. Purposes are hierarchies. Can be zero or more for each vocabulary.
Condition	Defines external context data (e.g., 'age', 'name', 'consentToMarketing', 'timeOfDay') that can be evaluated by conditions. Can be zero or more for each vocabulary.

Obligation	Defines actions that must be taken by the EPAL environment. Lists obligations (legislation and privacy policies) and/or additional steps that an enterprise is obligated to perform when a certain action occurs. For example, obligation "retention" with a parameter <days> 30</days> could denote that data shall be deleted after 30 days. Can be zero or more for each vocabulary.

Prior to EPAL's introduction, IBM researchers investigated enterprise privacy management. Karjoth *et al*. proposed a privacy policy model for enterprises to formalize privacy policies [KS02]. This model, identifies six privacy policy elements: subject, data, action, purpose, condition and obligation. The Authorization Specification Language [JSS01] is used to specify the authorization rules. Powers *et al*. are extending traditional access control mechanisms to enforce enterprise privacy policies and have proposed an enterprise privacy management architecture [PAS02, KSW02a, APS02].

Ashley *et al*. developed E-P3P (The Platform for Enterprise Privacy Practices) with which an organization can specify internal privacy practices in a way that can be enforced and thus guaranteed [AHK02, KSW02b]. Karjoth *et al*. have developed a method for automatically translating E-P3P policies into P3P policies, so that the organization can keep their privacy promises specified in P3P up-to-date with the actual practices [KSH03a]. This is an important step to ensure consistency between public P3P policies and internal privacy practices.

An essential concept in IBM's enterprise privacy management efforts is the "sticky policy paradigm" which means the policy that is specified and consented by data subjects at collection, and which governs data usage, holds true throughout the data's lifetime, even when the data is disclosed by one organization to another. Privacy policy specification is essential for enterprise privacy enforcement. Based on E-P3P, EPAL was developed as part of IBM's enterprise privacy management solution to specify enforceable enterprise privacy policies. After the EPAL specification's release, researchers began investigating its syntax and semantics. Backes *et al*. defined the formal syntax and semantics of EPAL policies and proposed an algorithm to efficiently compare EPAL policies [BBK04]. However, to date, no one has examined the expressiveness of EPAL in specifying natural language privacy polices, which motivates the work presented in this paper.

3. CIGNA Privacy Policy Case Study

Given our prior experience in extracting privacy goals from natural language policy documents and the clear need for standardization [AER01, AEB04], we sought to evaluate P3P and EPAL's ability to specify the practices expressed in these documents. Our study was intentionally exploratory, enabling us to familiarize ourselves with each specification language and evaluate their expressiveness as we now discuss.

3.1 Methodology

For this study, we analyzed two CIGNA privacy policy documents: (1) CIGNA HealthCare Notice of Privacy Practices [Cig04a] and (2) CIGNA's Public Online Privacy Statement [Cig04b]. The participating analysts experience levels were varied; four individuals comprised our team (1 Ph.D., 2 Ph.D. students and one B.S./M.S. student). Two analysts had previously conducted preliminary research on the EPAL specification language, and one analyst had done the same with P3P. All four analysts are very experienced with privacy policy analysis [AER01, AEB04]. Initially, we reviewed each language specification [AHK03, CDH04] to gain a general understanding of the syntax and requirements for the rules to be specified. This training enabled us to proceed with the policy specification exercise.

To record our analysis results, we set up a document in which we annotated the original natural language specifications with our analysis results. Our P3P statement specification effort entailed scanning both source policy documents, identifying all data-groups and creating statement blocks (using templates) for each data-group. Our EPAL rule specification process entailed creating a rule for each natural language policy statement that addressed a data access (again, using templates). For the P3P specification effort, we examined every sentence in each natural language policy document to identify the data categories mentioned in the document so that we could define the data-groups. We then elaborated a P3P statement template for each identified data category, data-ref, and data group. Table 3 portrays a sample P3P statement template in the two left hand columns, and the third column portrays the final P3P statement specification in P3P. In Table 3, the `data-group` is cookie information. For each data-group, a statement block was created with P3P elements: extension, consequence, purpose, recipient, and retention. Each of these elements is described in Table 1. In this example, cookie information secures online services, the recipient is ours, and retention is indefinite. Together all the elaborated statement block templates for each data-group constitute the entire policy schema, which includes other elements,

including dispute group, expiry, access etc.

Table 3: P3P Statement from CIGNA HealthCare Public Online Privacy Statement – Statement #3

Element Name	Element Value	P3P Rule
extension	N/A	<STATEMENT>
consequence	N/A	<EXTENSION/>
purpose	Other: to secure online services	<CONSEQUENCE>non-identifiable </CONSEQUENCE>
recipient	Ours	<PURPOSE><other>to secure online services </other></PURPOSE>
retention	Indefinitely	<RECIPIENT><ours/></RECIPIENT>
data-group	Cookie information	<RETENTION><indefinitely/></RETENTION>
		<DATA-GROUP>
		<DATA REF="#CIGNA.customer.cookieInfo"/>
		</DATA-GROUP>
		</STATEMENT>

For the EPAL specification effort, we first examined every sentence in each natural language policy document to identify data accesses. Each statement was examined by asking, "Does this statement refer to an information access?" If the answer to this question is "yes" we elaborated a new EPAL rule template, shown in the top half of Table 4.

Given *Statement A*, taken from the CIGNA HealthCare Notice of Privacy Practices, we derived the EPAL rule specification shown in Table 4.

> ***Statement A:*** *We may disclose your confidential information for research purposes, subject to strict legal restrictions.*

The bottom half of Table 4 shows the resulting EPAL rule specified in XML. For each rule, we specified the EPAL rule elements (ruling, user-category, action, data-category, purpose, condition, and obligation). In the following example the ruling is allow, user-category is CIGNA HealthCare, action is to disclose information, data-category is confidential information, purpose is for research and condition is that the disclosure is subject to strict legal restrictions. There is no obligation field in this particular example as that field is optional and there is no obligation in this statement.

Table 4: EPAL Rule Specification for CIGNA HealthCare Notice of Privacy Practices

Template Specification for Statement A EPAL Rule	
Paraphrased Statement	DISCLOSE confidential information for research purposes
ruling	ALLOW
userCategory	CIGNA Healthcare
action	DISCLOSE
dataCategory	Confidential Information
purpose	For research purposes
condition	Disclosure subject to strict legal restrictions
obligation	N/A

XML EPAL Rule for Statement A

```
<rule>
<ruling>ALLOW</ruling>
<userCategory>CIGNA Healthcare</userCategory>
<action>DISCLOSE</action>
<dataCategory>Confidential Information</dataCategory>
```

```
<purpose>Research Purposes</purpose>
<condition>Disclosure subject to strict legal restrictions</condition>
<obligation/>
</rule>
```

Note that we intentionally chose not to employ P3P or EPAL specification editors. This was to ensure that our evaluation of the specification languages was in no way biased by usability factors (by changing the focus of evaluation to the user interface rather than the underlying specification language). As for the artifacts produced, we specified six P3P statements for both of CIGNA's policy documents. In contrast, we specified 71 EPAL rules for both policies. In no way do we advocate a simplistic comparison of "apples and oranges." Instead, our intent is to factually report the resulting number of artifacts from our analysis using the two different specification languages.

3.2 Criteria for Evaluating Requirements and Policy Specification Artifacts

As previously mentioned, software requirements and privacy policies have much in common. Thus, given that our work is predicated on requirements engineering and requirements specification principles, we evaluated the P3P statements and EPAL rules that we produced as a result of our specification activities using common requirements engineering criteria, based upon the requirements "quality gateway" [RR99]. The "quality gateway" helps requirements engineers check the completeness and correctness of the requirements that they have collected from various sources before those requirements are added to the final requirements specification. If a requirement passes the quality test (i.e. passes through the gateway), it is added to the specification; if it does not pass the quality test, it must either undergo further elaboration or even elimination. The criteria that are commonly used to assess whether a requirement meets a minimal level of quality include: completeness, traceability, consistency, relevancy, correctness, ambiguity and creep. We now define quality gateway criteria as they are employed in the requirements specification arena and justify their relevance in privacy policy specification.

Completeness refers to the extent to which the elements of each requirement are comprehensively specified in terms that reflect everything that is known about that requirement for its given context. Requirements are often specified using templates and the elaborated templates are then evaluated against the quality gateway criteria. As we discuss below, we elaborated P3P statement templates and EPAL rule templates in the same way that we typically use requirements templates to specify goals, constraints, stakeholders, as well as pre- and post-conditions. Evaluating completeness in the context of policy specification seems appropriate given the similar use of templates in requirements and policy specification.

Traceability refers to the ability to track each requirement's source and origin. It is a measure of quality that reduces the risk of, for example, not propagating changes across lifecycle artifacts. Our studies have shown that organizational privacy policies are subject to frequent changes, either in response to newly introduced legislation or to reflect changes in an organizations' mission or business activities [AEB04]. Traceability helps prevent adverse effects as a result of improper change propagation in the context of both requirements and policy specification.

Consistency means that the specified requirements contain no contradictions with respect to either functionality or terminology. When software requirements are inconsistent, they are subject to misinterpretation and can result in software failures. When privacy policies are inconsistent, they are also subject to misinterpretation and can leave the respective organization legally vulnerable.

Relevancy refers to whether a particular potential requirement has legitimate bearing on the system at hand. Some requirements, even when well-stated and complete, are not relevant to the envisioned system; these requirements should be discarded. Most natural language privacy policy documents are long and much of the information they contain is not germane to the organizations' actual information practices. This criteria is especially useful for privacy policy specification as it helps analysts remain focused on that which is truly relevant for machine readable and/or machine enforceable policies.

Correctness refers to whether a requirement is consistent with other requirements, redundant or in conflict with another requirement, whether a requirement is testable, and unambiguous. Within the context of privacy policy specification, testability is an interesting concept because a policy that must be machine enforceable should certainly be testable. Similarly, a policy that must be machine-readable should also be testable.

Ambiguity in software requirements means that a particular requirement is subject to different interpretations and suggests that a particular requirement is not sufficiently refined to ensure testability.

Because privacy policy documents are laden with legalese and notoriously ambiguous and non-committal [AER01, AEB04], it is especially incumbent upon us to ensure that the policy rules and statements that we specify are unambiguous so as to ensure that our specifications are correct.

Creep refers to significant additions or modifications to a software system's requirements throughout the development lifecycle, resulting in extensions to and alteration of the software's functionality and scope. Requirements creep can be especially troublesome to developers due to the detrimental impact such changes may have on cost, resources, quality, or the ability to deliver a system that incorporates the new requirements on time. Organizational privacy policies are often written in a way that implies an expectation for policies to change in the future and sometimes express things that may happen in certain contexts in response to changes in the environment. As such, it is important to limit the ill effects of policy creep in our policy specifications.

In Section 4 we discuss whether the P3P and EPAL specification languages are rich enough to enable one to produce policy specifications that satisfy these criteria.

4. Lessons Learned

Throughout this case study, we gleaned valuable insights about the expressiveness of policy specification languages as well as the original natural language specification, as we now discuss.

P3P is a data-centric, whereas EPAL is an access-centric specification language.

The main difference between P3P and EPAL, from the specification perspective, is on the information focus of each language. Specifically, P3P statements are data-centric, whereas EPAL rules are access-centric. For example, a P3P statement derived from the CIGNA HealthCare policy documents described all the collection, use, storage, etc. of `CIGNA.customer.confInfo` (Confidential Customer Information). However, in EPAL there is a collection of rules that refer to specific information accesses/uses of Confidential Customer Information, and each individual rule is not all-inclusive of the practices related to that specific type of data. The way that EPAL specifications are written somewhat parallels the way that natural language policies are written, in that each derived EPAL rule came from an individual natural language policy statement. However, with P3P, each data-centric statement was derived from the privacy policy documents in their entirety. Thus, the P3P format does not relate to natural language policy structure as intuitively as EPAL does, so the process of deriving a P3P policy from a natural language policy is slightly different and less intuitive.

Privacy policies and specification languages must be able to express information disclosures recipients.

End users are typically concerned about who will collect, receive, and store their personal information. Unfortunately, studies have shown that there exists a mismatch between what privacy policy documents express and what end users want these documents to express [EAA03]. In CIGNA's natural language policies, we observed cases in which insufficient information was provided about information recipients in the event of a disclosure. For example, the CIGNA Notice of Privacy Practices states: "We may use or disclose your confidential information to provide you with a promotional gift of nominal value." Admittedly, the natural language policy is rather ambiguous about the possible data recipients to begin with. Nonetheless, when we attempted to express this policy as an EPAL rule, it became evident that EPAL allows one to express that a possible disclosure may occur. However, it does not allow one to express to whom the disclosure is made. As shown the example below (EPAL rule 12.b), we specified the EPAL `user-category` as the organization disclosing the information (CIGNA), but in reality the user category should be the organization that is using the data, which in this case may or may not be CIGNA.

```
12.b)
DISCLOSE confidential info. to provide you with a promotional gift of nominal value
    ruling:                 ALLOW
    userCategory:           CIGNA healthcare
    action:                 DISCLOSE
    dataCategory:           Confidential Information
    purpose:                to provide you with a promotional gift of nominal value
    condition:              n/a
```

End-users may view the fact that the policy does not specify the possible recipients of this information disclosure as a vulnerability. In some cases, the natural language policy does state the recipient; however, a work-around is needed to enable one to document this recipient information. In the EPAL rule below (8.b), we noted the recipient in the <condition> field.

```
8.b) DISCLOSE confidential information to audit employee health benefit plan
    ruling:                 ALLOW
    userCategory:           CIGNA HealthCare
    action:                 DISCLOSE
    dataCategory:           Confidential Information
    purpose:                audit employee health benefit plan
    condition:              recipient must be customer's employer or company acting on
                            employer's behalf
```

In contrast to EPAL, P3P does allow one to specify whether or not there is a recipient. However, the specification is open to multiple interpretations because `<other-recipient>` is not as explicit as one would desire as shown in the P3P statement derived from *Statement B* in the CIGNA HealthCare Notice of Privacy Practices.

> ***Statement B:*** *We also may disclose your confidential information to another health plan or a provider who has a relationship with you, so that it can conduct quality assessment and improvement activities - for example, to perform case management.*

```
<STATEMENT>
       <PURPOSE>
              <other-purpose required="always">
              for 3rd party quality assessment and improvement
              </other-purpose>
       </PURPOSE>
<RECIPIENT><other-recipient/></RECIPIENT>
<RETENTION><indefinitely/></RETENTION>
       <DATA-GROUP>
              <DATA REF="#CIGNA.customer.confInfo"/>
       </DATA-GROUP>
</STATEMENT>
```

Because information recipients are relevant to information systems and transactions, our study revealed that P3P helped us ensure that relevant recipient data was expressed in our P3P statements. In contrast, we were unable to ensure this kind of relevancy in our EPAL rules.

The ability to specify multiple responsible organizations in EPAL is necessary for the language to more fully support the "Sticky Policy Paradigm."

A major challenge facing privacy policy researchers is how to help organizations manage large data sets with heterogeneous privacy practice requirements. These situations arise when data is collected over a period of time during which an organization's privacy policies are changing. Consider a situation in which data is collected before and after a policy change that affects that data. One piece of data (e.g. name) may be protected from disclosure, whereas another piece of similar data (e.g. another name collected after a policy change) may not be protected. There must be some way to track the policies that apply to a specific piece of data. The "Sticky Policy Paradigm" seeks to address this by ensuring that the applicable policy for a piece of data is "stuck" to that data and follows it for its entire lifetime [KSW02b]. As previously mentioned, the sticky policy paradigm allows policies to be specified and consented to by the data subjects at the time of collection. Those policies then govern data usage during the data's lifetime even after the data is disclosed by one organization to another.

Our case study revealed that, in its current form, EPAL does not adequately support the sticky policy paradigm because it provides no way to specify multiple responsible organizations in the policy when more than one organization is granted access to the same piece of data. For example, suppose company A collects personal data from customers. At the time of collection, company A is the responsible organization for the EPAL policy associated with this data. In company A's EPAL policy, we can specify who has issued the policy, information about them, the policy's effective date, etc. Later this data is disclosed to company B, which becomes the responsible organization for this data. Company B is supposed to enforce the same policy according to the Sticky Policy Paradigm. However, there is currently no direct way to specify information about company B in the original EPAL policy. As an example, consider *Statement C* from the CIGNA HealthCare Notice of Privacy Practices, the information disclosure recipients, "your employer or a company acting on your employer's behalf", are additional organizations

responsible for maintaining information confidentiality.

> ***Statement C:*** *"We may disclose your confidential information to your employer or to a company acting on your employer's behalf, so that it can monitor, audit and otherwise administer the employee health benefit plan in which you participate."*

Unfortunately, we cannot specify information disclosure recipients in `userCategory` because the recipients in this statement are not CIGNA actors and their behavior is not enforceable by CIGNA's EPAL policy. Thus, as a workaround, we specified the information disclosure recipient as a condition in this rule:

```
8.a) DISCLOSE confidential information to monitor employee health benefit plan
   ruling:                ALLOW
   userCategory:          CIGNA HealthCare
   action:                DISCLOSE
   dataCategory:          Confidential Information
   purpose:               monitor employee health benefit plan
   condition:             recipient must be customer's employer or company acting on
                          employer's behalf.
```

Although we created this temporary workaround for the problem by specifying the responsible organization within the `<condition>` tag as shown in the example above, this is admittedly a less-than-optimal solution. Thus, it would be helpful for EPAL to allow this sort of specification in its policies.

Both P3P and EPAL are insufficient for specifying high-level company obligations.

P3P and EPAL are both designed to specify an organization's privacy practices as they relate to a specific website or transactional system, respectively. Neither language is meant to express an organization's overall privacy practices. Instead, they are intended to express the specific subset of organizational privacy practices that is relevant to the language's purpose. For example, neither P3P nor EPAL can readily express *Statement D* below from CIGNA's HealthCare Notice of Privacy Practices:

> ***Statement D:*** *CIGNA HealthCare locations that maintain confidential information have procedures for accessing, labeling and storing confidential records. Access to our facilities is limited to authorized personnel.*

Statement D expresses high-level company obligations, but the scope of the P3P and EPAL languages lack the breadth necessary to encompass the whole organization and all of its information handling practices. Natural language policies, however, typically cover such statements [AER01]. To this end, neither P3P nor EPAL truly helped us ensure that our resulting policy specifications were complete, as advocated quality criteria.

P3P cannot express "Effective On" dates for its policies

Effective dates are important from the end-user's perspective as they seek to better manage their personal information. If an end-user provided information in the past, and are considering providing more, they may wish to know how long the current privacy policies have been in effect and governing the organization's information handling practices. P3P provides an `<EXPIRY>` tag that lets end-users, or their agents, determine how long a policy will be effective, but there is no way to specify a time or date upon which the policy went into effect. Natural language policies frequently contain statements such as "This Notice is effective on April 14, 2003" (as found in CIGNA's HealthCare Notice of Privacy Practices). This is important information that cannot be expressed in P3P. In this respect, the language does not allow policy makers to ensure that their policy specifications are complete, as advocated by our quality criteria. However, a simple solution to this dilemma is to add an `<effective-on>` tag to the specification language.

P3P has limited scope as a public privacy policy specification language.

Any policy specification language will have some predefined scope that it was created to cover. EPAL, for example, is meant to express internal privacy practices in a machine-readable and enforceable way. In contrast, P3P is intended to express external privacy promises in a machine-readable fashion that user agents can parse and evaluate during a user's Internet browsing. P3P's scope only covers a specific, limited subset of external privacy promises. More specifically, P3P can only express privacy promises related to specific information collection instances on an organization's website. It cannot express the general privacy policies of the organization as a whole. For example, consider *Statement E*, which was taken from the CIGNA HealthCare Notice of Privacy Practices:

> ***Statement E:*** *CIGNA HealthCare locations that maintain confidential information have procedures for accessing, labeling and storing confidential records. Access to our facilities is limited to authorized personnel.*

This policy statement contains information regarding information protection practices that exist outside the scope of the organization's website (i.e. it describes some of the organization's physical security infrastructure and procedures). Such an information handling policy cannot be expressed in P3P, as it is not associated with a specific information instance from the website. We believe P3P's scope is too narrow and could be improved by incorporating the ability to express such statements with the language, thus allowing policy makers to fully disclose an organization's privacy practices.

EPAL and P3P specification can help reduce ambiguities and omissions that appear in natural language policy documents.

As we have observed in previous studies [AER01, AEB04, AHB04], website natural language privacy policy documents are typically laden with inconsistencies. These policies are also often ambiguous and contain statements that conflict with each other [AEB04], increasing the likelihood that the policy statements they contain will be misinterpreted. In our previous studies, we employed a goal-driven technique that allows one to examine policy documents at a fine level of granularity – at the individual statement level. Whereas the goal-driven approach proved effective in identifying ambiguities, conflicts, and omissions, we found the use of P3P and EPAL much more efficient.

As an example of an ambiguity discovered, consider *Statement F* from CIGNA's HealthCare Notice of Privacy Practices:

> ***Statement F:*** *CIGNA HealthCare also discloses confidential information to accreditation organizations such as the National Committee for Quality Assurance (NCQA) when the NCQA auditors collect Health Plan Employer Data and Information Set (HEDIS®)** data for quality measurement purposes. When we enter into these types of arrangements, we obtain a written agreement to protect your confidential information.*

One of our analysts interpreted the written agreements as being between CIGNA and individual customers, while another analyst interpreted the written agreement as being between CIGNA and NCQA auditors. Clearly, multiple interpretations are possible here. However, our analysis activities, that required us to think about each policy statements from both a data-centric perspective and an access-centric perspective, enabled us to develop a deeper understanding of such statements because the structured analysis forced us to resolve the ambiguity to clearly specify the policy using our templates. Thus, we found that the process of analyzing the policy documents according to what EPAL and P3P require, respectively, helped us resolve these ambiguities. For *Statement E*, our analysis team unanimously reached an understanding that the original statement was referring to agreements between CIGNA and NCQA auditors, resulting in EPAL rule 11.a:

```
11.a) DISCLOSE confidential information for quality measurement
    ruling:                 ALLOW
    userCategory:           CIGNA healthcare
    action:                 DISCLOSE
    dataCategory:           Confidential Information
    purpose:                for quality measurement
    condition:              written agreement from NCQA required
```

Once *Statement E* was specified in EPAL, its meaning was clear and the policy could no longer be misinterpreted. This is a clear benefit of semi-structured policy specification languages over their natural language counterparts. Moreover, our use of both specification languages helped us ensure that our resulting specifications were not ambiguous — one of our quality criteria.

Our analysis also helped us discover omissions in the natural language policy. For example, the following statements appeared in the CIGNA HealthCare Notice of Privacy Practices:

> ***Statement G:*** *We may disclose your confidential information to your employer or to a company acting on your employer's behalf, so that it can monitor, audit and otherwise administer the employee health benefit plan in which you participate.*

and

> ***Statement H:*** *Your employer is not permitted to use the confidential information we*

disclose for any purpose other than administration of your health benefit plan.

A loop hole was identified in *Statement G* because it states that that a customer's employer is not allowed to use the confidential information in a certain way, however, there is no statement in the natural language policy that covers "an agent working on the employer's behalf." Because *Statement F* states that information may be disclosed to this third party, it is a clear omission in the natural language policy that there is no information regarding that organization's treatment of the disclosed data.

5. Summary and Future Work

It is increasingly important for organizations to construct accurate and effective privacy policies that document their information handling and usage practices. In this paper, we show how quality criteria used in software requirements specification can be used to evaluate privacy policies specified using P3P and EPAL. We provided criteria and evaluated P3P and EPAL to examine how well they produce specification artifacts that satisfy these criteria, as we now discuss.

Our discussion was framed by our consideration of RE quality criteria and how well the respective languages enabled us to produce quality policy specifications. EPAL allowed us to produce more complete rules than P3P because we were able to complete all required tags in our EPAL rule templates. In contrast, were unable to complete all our P3P statements in the same manner. However, as previously discussed, EPAL does not have a `<RECIPIENT>` tag. Thus, although we were able to evaluate the languages based upon this criteria, the results do not necessarily reflect all the areas where a specific language may be lacking. Additionally, it is very important to be able to trace each policy element in a specification to its source. Being XML-based, both specification languages can support arbitrary attribute definitions that could be used for this purpose, but neither the P3P nor EPAL's language specification mentions this functionality. We advocate adding this to each specification, respectively. Moreover, this functionality should be made more apparent to end-users.

Both languages partially support the consistency criteria, as they both allow for a standardized vocabulary, but if this feature is used incorrectly, then there remains room for misinterpretation because there is no limit to the vocabulary terms that could be used. Within the context of our study, the relevancy criteria is especially interesting. Not all the statements in the natural language policy are relevant, but transforming the natural language policy statements into P3P statements and EPAL rules helped us weed out the irrelevant statements, and only the relevant ones made it into the final policy documents. Thus the specification process in both the P3P and EPAL efforts helped ensure that the relevancy criteria was satisfied in the final policy documents.

EPAL only partially satisfies the correctness criteria because there are many statements that we could not operationalize given only the information in the natural language policy document. Because P3P is not enforceable and because it is subject to different interpretations by different user agents [LYA03], it was deemed insufficient in this regard, and we view the resulting P3P statements as inherently ambiguous. In contrast, EPAL provides support for vocabularies to reduce terminology ambiguities. During this study, we were unable to properly evaluate the Creep criteria because proper evaluation would require us to evaluate two versions of the same document that changed over a period of time.

Our study is limited in that the only source to specify P3P and EPAL policies was the natural language privacy policies. Because we are not CIGNA employees, we were unable to obtain additional information (e.g. via interviews) that would have clarified ambiguities. Additionally, we did not have detailed information about the organization's existing privacy practices and business processes; thus, we were unable to operationalize many of the EPAL policy rules that were too vague for enforcement. Even though the study is limited, our findings are insightful. Every privacy policy in an organization, whether expressed in P3P, EPAL, or some other standard, must be in compliance with the organization's natural language privacy policies. The natural language policies serve as an organization's public promise to its customers. Our study revealed that many promises expressed in natural language privacy policies are neither expressible in P3P nor enforceable with EPAL. Our novel use of the requirements engineering quality criteria as a framework for our discussion and evaluation of the features of both languages allowed us to suggest ways in which both languages can be extended.

Our plans for future work entail a more extensive study involving a more comprehensive specification effort. We also plan to compare policy specifications from natural language policies with specifications derived from goals. We are developing a Security and Privacy Requirements Analysis Tool (SPRAT) [JAS04] to support this kind of analysis and as well as policy and requirements specification.

Acknowledgements

The authors thank Calvin Powers, Jack Frink, Bharathy Sethumadhavan and Matthew Vail for their comments about this paper as well as several individuals who answered specific questions about the P3P and EPAL specifications: Paul Ashley, Lorrie Cranor and Matthias Schunter. This work was funded by National Science Foundation ITR Grant #0325269 and the CRA Distributed Mentor Program.

References

[AEB04] A.I. Antón, J.B. Earp, D. Bolchini, Q. He, C. Jensen and W. Stufflebeam. The Lack of Clarity in Financial Privacy Policies and the Need for Standardization, *IEEE Security & Privacy*, 2(2), pp. 36-45, 2004.

[AEC03] A.I. Antón, J.B. Earp and R.A. Carter. Precluding Incongruous Behavior by Aligning Software Requirements with Security and Privacy Policies, *Information and Software Technology*, Elsevier, 45(14), pp. 967-977, 1 November 2003.

[AEP01] A.I. Antón, J.B. Earp, C. Potts and T.A. Alspaugh. The Role of Policy and Privacy Values in Requirements Engineering, *5th IEEE International Symposium on Requirements Engineering (RE'01)*, Toronto, Canada, pp. 138-145, 27-31 August 2001.

[AER01] A.I. Antón, J.B. Earp and A. Reese, Goal Mining to Examine Health Care Privacy Policies, NCSU Technical Report TR-2001-10, 6 November 2001.

[AHB04] A.I. Antón, Q. He, and D. Baumer. The Complexity Underlying JetBlue's Privacy Policy Violations, To Appear: *IEEE Security & Privacy*, 2004

[AHK02] P. Ashley, S. Hada, G. Karjoth and M. Schunter. E-P3P Privacy Policies and Privacy Authorization. *Proc. of the Workshop on Privacy in the Electronic Society (WPES'02)*. Washington D.C. November 21, 2001.

[AHK03] P. Ashley, S. Hada, G. Karjoth, C. Powers and M. Schunter. Enterprise Privacy Authorization Language (EPAL 1.1) Specification. IBM Research Report. http://www.zurich.ibm.com/security/enterprise-privacy/epal. 2003.

[APS02] P. Ashley, M. Schunter, and C. Powers. From Privacy Promises to Privacy Management: A New Approach for Enforcing Privacy Throughout an Enterprise, *Proc. of the ACM New Security Paradigms Workshop*, 2002.

[BBK04] M. Backes, W. Bagga, G. Karjoth, and M. Schunter. Efficient Comparison of Enterprise Privacy Policies, *Proc. of the 19th ACM Symposium on Applied Computing (SAC'04)*, pp. 375-382, 2004.

[BCK03] S. Byers, L.F. Cranor, D. Kormann," Automated analysis of P3P-enabled Web sites," Proceedings of the 5th international conference on Electronic commerce, Pennsylvania, pp 326-338. September 2003.

[BHA04] D. Bolchini, Q. He, A.I. Antón and W. Stufflebeam. I need it now: Improving Website Usability By Contextualizing Privacy Policies, To appear: The 4th International Conference on Web Engineering (ICWE 2004), Munich, Germany, 28-30 July 2004.

[CDH04] L. Cranor, B. Dobbs, G. Hogben, J. Humphrey, M. Langheinrich, M. Marchiori, M. Presler-Marshall, J. Reagle, M. Schunter, D.A. Stampley, R. Wenning. *The Platform for Privacy Preferences 1.1 (P3P1.1) Specification*, http://www.w3.org/TR/P3P11. W3C Working Draft 27 April 2004.

[CIN03] R. Crook, D. Ince, and B. Nuseibeh. Modelling Access Policies Using Roles in Requirements Engineering, *Information and Software Technology*, 45 (14), pp. 979-991, Elsevier, 2003.

[Cig04a] CIGNA HealthCare. Public Online Privacy Statement. http://www.cigna.com/general/privacy/public.html. Accessed June 2004.

[Cig04b] CIGNA HealthCare. Notice of Privacy Practices. http://www.cigna.com/general/privacy/healthcare/standard.html. Accessed June 2004

[Cra03] L. Cranor. P3P: Making privacy policies more useful. *IEEE Security & Privacy*. 1(6), pp. 50–55, November-December 2003.

[Dam02] N.C. Damianou. A Policy Framework for Management of Distributed Systems, *PhD Thesis*, Imperial College, London, 2002.

[DLF93] A. Dardenne, A. van Lamsweerde, and S. Fickas. Goal-Directed Requirements Acquisition, *Science of Computer Programming*, 20: pp. 3-50, 1993.

[EAA03] J.B. Earp, A.I. Antón, L. Aiman-Smith, and W.H. Stufflebeam. Crossed Signals: Internet Privacy Policies and User Concerns, Submitted to: *IEEE Transactions on Engineering Management,* 2003.

[EM99] D. Eames and J.D. Moffett. The Integration of Safety and Security Requirements, *Proceedings of the 18th International Conference on Computer Safety, Reliability and Security*. pp.468-480. 27-29 September 1999.

[Fon01] P.J. Fontaine. Goal-Oriented Elaboration of Security Requirements, Project Dissertation, Université Catholique de Louvain, Belgium, 2001.

[FTC98] Federal Trade Commission. Privacy online: A report to congress, http://www.ftc.gov/reports/privacy3/. June 1998.

[Hoc02] H. Hochheiser, The platform for privacy preference as a social protocol: An examination within the U.S. policy context, *ACM Transactions on Internet Technology*, 2(4), pp 276-306, November 2002.

[Jam04] L. Jamtgaard and The Internet Education Foundation. The P3P Implementation Guide. As of 5-5-2004, http://p3ptoolbox.org/guide/.

[JAS04] N. Jain, A.I. Antón, W.H. Stufflebeam and Q. He. *Security and Privacy Requirements Analysis Tool (SPRAT) Software Requirements Specification*, NCSU CSC Technical Report TR-2004-7, February 24, 2004.

[JSS01] S. Jajodia, P. Samarati, M.L. Sapino and V.S. Subrahmanian. Flexible Support for Multiple Access Control Policies, *ACM Transactions on Database Systems*, 26(2), pp. 214-260, 2001.

[KS02] G. Karjoth and M. Schunter. A Privacy Policy Model for Enterprises, *Proc. of the 15th IEEE Computer Security Foundations Workshop*, pp. 271-281, 2002.

[KSW02a] G. Karjoth, M. Schunter, and M. Waidner. Privacy-enabled Services or Enterprises. Proc. of the 13th IEEE International Workshop on Database and Expert Systems Applications (DEXA'02), 2002.

[KSW02b] G. Karjoth, M. Schunter, and M. Waidner. The Platform For Enterprise Privcy Practices ^ Privacy-enabled Management of Customer Data. Proc. of the 2002 Workshop on Privacy Enhancing Technologies, 2002.

[LYM03] L. Liu, E. Yu and J. Mylopoulos. Security and Privacy Requirements Analysis within a Social Setting, *Proc. of the 11th International Requirements Engineering Conference (RE'03)*, pp. 151-161, 2003.

[Mof99] J.D. Moffett. Requirements and Policies, Workshop on Policies in Distributed Systems, HP-Laboratories, Bristol, UK, 1999.

[NE00] B.A. Nuseibeh and S.M. Easterbrook. Requirements Engineering: A Roadmap, In A.C.W. Finkelstein (ed.) *The Future of Software Engineering*. (Companion Volume to the *Proc. of the 22nd International Conference on Software Engineering, ICSE'00*). IEEE Computer Society Press, 2000.

[PAS02] C. Powers, P. Ashley, and M. Schunter. Privacy promises, access control, and privacy management. enforcing privacy throughout an enterprise by extending access control. In Proceedings of the Third International Symposium on Electronic Commerce, pp. 13–21, October 2002.

[Pfl91] S.L. Pfleeger. A Framework for Security Requirements, Computers and Security. 10(6), pp. 515-523, 1991.

[Pre02] M. Presler-Marshall. The platform for privacy preferences 1.0 deployment guide. February 2002, http://www.w3.org/TR/2002/NOTE-p3pdeployment-20020211.

[RR99] S. Robertson and J. Robertson. *Mastering the Requirements Process*. Addison-Wesley. New York, 1999.

[Sch03] M. Schunter. Enterprise Privacy Authorization Language (EPAL 1.1). http://www.zurich.ibm.com/security/enterpriseprivacy/epal/Specification/index.html, 2003.

[SHW04] M. Schunter, E. V. Herreweghen, and M. Waidner. Translating epal to p3p - how to keep enterprise privacy promises in sync with the acutal practices. As Of 5-5-2004, http://www.w3.org/2003/p3p-ws/pp/ibm2.html.

[Yu93] E. Yu. Modeling Organizations for Information Systems Requirements Engineering, *Proc. of the 1st IEEE International Symposium on Requirements Engineering*, pp. 34-41, 1993.