



NRC Publications Archive Archives des publications du CNRC

Privacy Management Architectures for E-Services Korba, Larry; Song, Ronggong; Yee, George

This publication could be one of several versions: author's original, accepted manuscript or the publisher's version. /
La version de cette publication peut être l'une des suivantes : la version prépublication de l'auteur, la version acceptée du manuscrit ou la version de l'éditeur.

NRC Publications Record / Notice d'Archives des publications de CNRC:
<https://nrc-publications.canada.ca/eng/view/object/?id=a1181942-b291-4b5f-822e-fe36a9664a68>
<https://publications-cnrc.canada.ca/fra/voir/objet/?id=a1181942-b291-4b5f-822e-fe36a9664a68>

Access and use of this website and the material on it are subject to the Terms and Conditions set forth at
<https://nrc-publications.canada.ca/eng/copyright>
READ THESE TERMS AND CONDITIONS CAREFULLY BEFORE USING THIS WEBSITE.

L'accès à ce site Web et l'utilisation de son contenu sont assujettis aux conditions présentées dans le site
<https://publications-cnrc.canada.ca/fra/droits>
LISEZ CES CONDITIONS ATTENTIVEMENT AVANT D'UTILISER CE SITE WEB.

Questions? Contact the NRC Publications Archive team at
PublicationsArchive-ArchivesPublications@nrc-cnrc.gc.ca. If you wish to email the authors directly, please see the first page of the publication for their contact information.

Vous avez des questions? Nous pouvons vous aider. Pour communiquer directement avec un auteur, consultez la première page de la revue dans laquelle son article a été publié afin de trouver ses coordonnées. Si vous n'arrivez pas à les repérer, communiquez avec nous à PublicationsArchive-ArchivesPublications@nrc-cnrc.gc.ca.





National Research
Council Canada

Conseil national
de recherches Canada

Institute for
Information Technology

Institut de technologie
de l'information

NRC - CNRC

Privacy Management Architectures for E-Services *

Korba, L., Song, R., and Yee, G.
2006

* published in Privacy Protection for E-Services, published by Idea Group
Inc. 2006. NRC 48271. Yee, G. (Editor)

Copyright 2006 by
National Research Council of Canada

Permission is granted to quote short excerpts and to reproduce figures and tables
from this report, provided that the source of such material is fully acknowledged.

Privacy Management Architectures for E-Services¹

Larry Korba

National Research Council Canada
Institute for Information Technology
1200 Montreal Road, Building M-50
Ottawa, Ontario, Canada K1A 0R6
Phone: 613- 998-3967
Fax: 613 952-7151
Email: Larry.Korba@nrc.ca

Ronggong Song

National Research Council Canada
Institute for Information Technology
1200 Montreal Road, Building M-50
Ottawa, Ontario, Canada K1A 0R6
Phone: 613-990-6869
Fax: 613 952-7151
Email: Ronggong.Song@nrc.ca

George Yee

National Research Council Canada
Institute for Information Technology
1200 Montreal Road, Building M-50
Ottawa, Ontario, Canada K1A 0R6
Phone: 613-990-4284
Fax: 613 952-7151
Email: George.Yee@nrc.ca

Privacy Management Architectures for E-Services¹

ABSTRACT

There have been a number of recent developments in architectures for privacy management. These architectures may be applied to the development of e-services. This chapter describes some driving forces and approaches for the development and deployment of a privacy architecture for e-services and reviews several architectures that have been proposed or developed for managing privacy. The chapter offers the reader a quick tour of ideas and building blocks for creating privacy-protection enabled e-services and describes several privacy information flow scenarios that can be applied in assessing any e-service privacy architecture. The chapter concludes with a summary of the work covered and a discussion of some outstanding issues in the application of privacy architectures to e-services.

KEYWORDS: privacy, privacy management, privacy protection, service, e-service, web service, architecture

INTRODUCTION

Before describing several different architectures for managing privacy, it is worthwhile to describe briefly the privacy and e-services landscape. This section outlines the context and general approach for privacy architecture development.

Background and Context

Over the past 6 years, major companies have used web services (i.e. Internet-enabled services) and e-services (network enabled services) interchangeably. For the purposes of this chapter we will use the term e-service to apply to either a web service (non-standards based Internet-enabled service) or a Web Service (XML standards based Internet-enabled service). E-services mean different things to technical people and business people. From the business context, e-services are described as an emerging paradigm that offers increased efficiency, enhanced services and stronger customer relationships through Internet-enabled applications that are reusable and customizable to user needs. E-services may be applied to Business-to-Consumer, or Business-to-Business situations. Moreover, the approach with e-services is to provide more value to customers. Adding value involves discerning what clients want. A service supplier may attempt to discern wants and needs through questionnaires or surveys, inferences from other data sources, or through direct requests from the consumer.


From the technical point of view, standards based e-services refer to a set of programming standards that makes the interplay between different types of software over the Internet happen without human intervention. These standards include Extensible Markup Language (XML), Standard Object Access Protocol (SOAP), Web Services Description Language (WSDL) and a variety of other web services definition languages. Middleware is built around these standards to support delivering technology to a customer over the Internet. For the purposes of this chapter, we can simply define an e-service as: a service or resource made available on the Internet.

The value service providers offer to customers in recent times stem from increasingly personalized services. Personalized services are selected on the basis of the needs and desires of clients. These are often directly associated with the name, and other personally identifiable information associated with the customer. In fact, in order to determine possible follow-on services in which a client may be interested, a service provider may resort to data mining from many different sources, collecting or inferring information about a client that may be quite personal. Considering the acceleration in technology development in support of deploying new services, the growing variety of services being developed, and the underlying approach of compiling, storing and analyzing information about users in an attempt to increase service value, it is clear that there are significant pressures on privacy. The pressures to build service applications rapidly to meet the new revenue opportunities also lead to questions regarding the implementation of security technology in support of privacy functions.

It is important to understand that the concept of privacy from the legal perspective is in disarray (Soslove, 2002). Without a consistent definition of privacy, adjudication and law-making do not fare well against the concrete and competing interests of other parties. Similarly, attempting to build privacy technologies based upon legal compliance is tantamount to building a product without appropriate requirements specifications. In fact, researchers and developers today often base their concepts and developments on core ideals related to privacy. Privacy principles, for instance, such as those compiled by the Canadian Standards Association (CSA) provide some general guidelines that have been used to form the basis of some technology developments.

Yet there are several aspects other than legislation that lead to determining technology and procedures to be put in place for e-service privacy. At this point it is worthwhile to examine the contributing factors to the need for privacy from the organizational perspective. Shaping users' or citizens' attitudes toward privacy are four items (Figure 1):

Figure 1. Citizens' attitudes towards privacy attitudes stem from 4 driving forces: Corporate Policy, Legislation, Social Norms, and Technology.



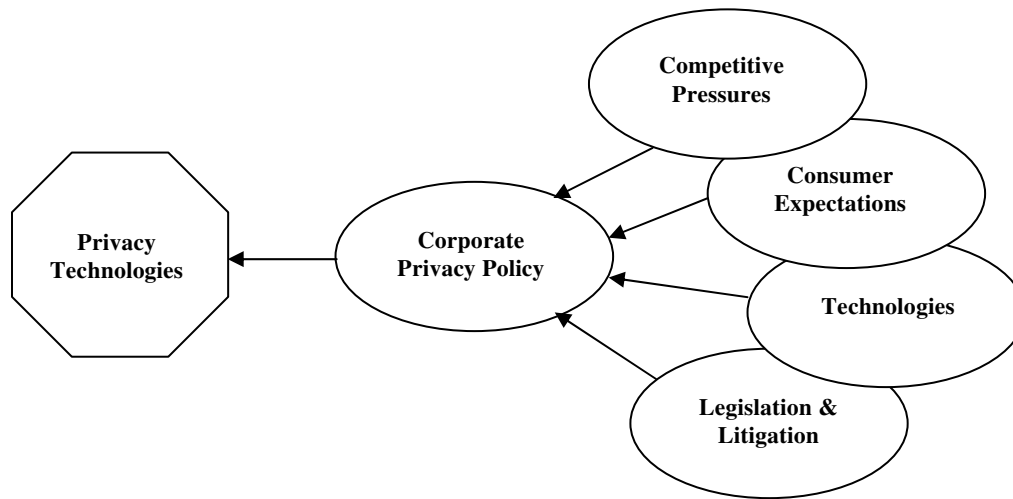
- Most organizations have organizational policies for dealing with personally identifiable information. In this case, organizations may include government, private sector and not-for-profit corporations. The policy may be based upon the organization's philosophy towards its business or clients. The policy may also reflect legal requirements, requirements based upon the organization's business model or requirements of its partners.

- Legislation also affects privacy attitudes. If privacy is held highly for a country, its laws and the emphasis on compliance to those laws will be of a high standard and will influence organizational policies and social norms as well as overall privacy attitudes.
- Social Norms add another dimension to privacy attitudes. Social norms are modulated by circumstances. For instance, when users are not identified they tend to be freer with their personal information (Cranor, 1999).
- By technology we mean all sorts of technology in our environment. Technology affects privacy attitudes. In combination with where and how it is deployed or used, technology has an effect on privacy attitudes. For instance, cameras located in public areas along with notices of camera surveillance may lead to a reduction in shoplifting. However, a camera connected to a home computer being used by someone who is at ease with the technology may lead to exhibitionist behavior!

Clearly, there is interplay between all of the factors that drive privacy attitudes. Social Norms, for instance drive legislation, corporate policy, and technology development. Technology may be regulated by legislation, metered by organizational policy and driven by social norms. Context has another role to play in the perception of what is private (Cranor et al., 1999) and may be reasoned about in different ways (Lessig, 1998).

Within an organization, privacy technologies are developed and deployed based upon corporate policy which is developed and modulated by competitive pressures, consumer expectations, technology and legislation, and the result of any litigation regarding how other organizations have dealt with private data (Figure 2).

Figure 2. Corporate drivers for privacy technologies



- Technologies in Figure 2 refer to all technologies, but in particular those that may have a threat upon privacy of personal data (e.g. insecure databases, insecure protocols that may be used for an e-service, scanning and sniffing software, radio frequency identification tags, etc.).
- Competitive pressures arise mainly due to influences in a business environment, wherein services are offered at increasingly competitive prices, or at the same and/or lower prices, but with increased service levels.
- Consumer expectations drive what an e-service may offer as well as how the service provider is expected to deal with personal data.
- Service providers may be covered under certain legislation regarding privacy aspects of the services they offer and the nature of the data involved. Litigation and settlements with respect to privacy-related disputes drive corporate privacy policy. Procedures prescribing how data must be dealt with may be driven by pertinent legislation and the threat of litigation.

An organization’s technology requirements for maintaining its privacy policy may vary greatly from service to service. While there may be potential legal requirements for implementing privacy approaches, organizations take several steps to be in a position to develop and implement privacy enhanced processes. These steps are described in table 1.

Table 1. Privacy-sensitive approach towards service implementation. Note that this list does not cover all aspects of a service, only those pertaining to privacy

Step	Description/Notes	Involvement
1.Determine organizational roles and responsibilities	For the service to be provided, determine organizational roles and responsibilities	Legal, executive management, government, privacy officer
2.Develop organizational privacy policy and security policies	The policies will evolve over time, with changes in the organization, services provided, feedback from customers, legislative and technology changes.	Legal, management, IT staff, privacy officer
3.Educate/inform staff and outsiders	This step would have started with the results of step 1. Within the organization, at all levels, via written, electronic and oral communication inform and educate the staff about the roles of those responsible for privacy and security and the security and privacy policies themselves.	All staff.
4.Thoroughly understand the service to be deployed, especially regarding data collection, processing and storage.	Data collected: <ul style="list-style-type: none"> • Determining the data to be collected. Does it relate or is it linked with the identity of people using the service? In this analysis the objective is to minimize the amount of data collected, thus minimizing potential privacy exposures. • May the user select what is to be collected? • Why is the data needed? Some data must be collected to provide the service, other data may be 	Service architect, consultation with legal, management, privacy officer

	<p>used for logging or tracking purposes.</p> <ul style="list-style-type: none"> • How long is the data required? • Will there be a pseudonymous or anonymous service? • How and where will personal data be stored? <p>What sort of logging will the service application use and is it possible to discern personally identifiable information from log entries?</p>	
5. Develop a privacy policy for the service	Based upon organizational policy, legal or regulatory requirement, and requirements for the service and business arrangements related to the service, develop a privacy policy	Privacy officer, service architect, legal, management
6. Technical design and implementation of the service	<p>Items to consider:</p> <ul style="list-style-type: none"> • Privacy policy disclosure: How will it be revealed to the user? • User interfaces for the service need to be carefully designed and tested to build the trust levels of the user • Test it well 	Service architect, programming staff
7. Privacy impact assessment	It is advisable to have a qualified external professional perform a privacy impact assessment, because of the experience and objectivity such a company brings. It is worthwhile having an assessment done at early stages of the design to lessen the chance of privacy-impaired design.	Privacy officer, consulting professional, or internal staff with appropriate experience and latitude to perform the assessment.
8. Launch service	Monitor feedback from clients regarding privacy issues. Correct /improve procedures and/or the application to deal effectively with feedback.	Service architect, programming staff, privacy officer
9. Improve service	Based upon feedback from clients or technical assessments made of the service application and support networks or hardware. Steps 7, 8, and 9 may iterate representing new levels of services or improvements coming on stream.	Architect, programming staff, privacy officer, other IT staff, management, marketing

PRIVACY ARCHITECTURES FOR E-SERVICES

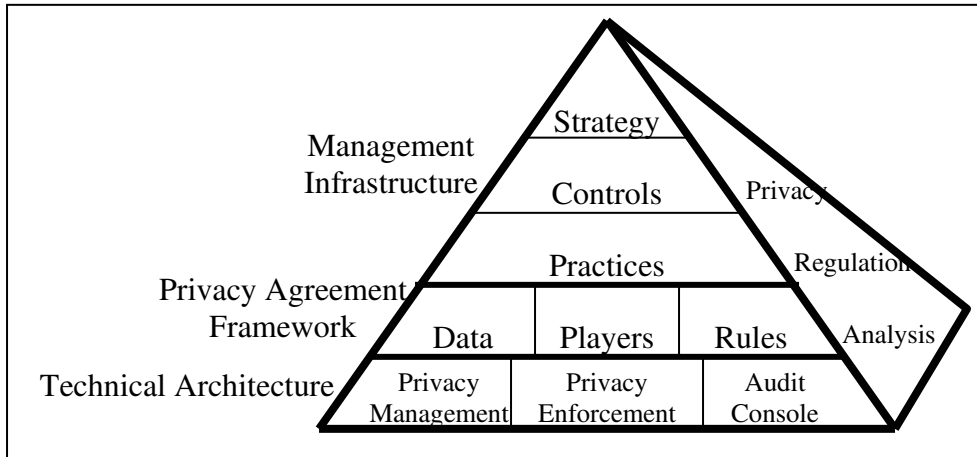
Having covered background information regarding privacy and e-services and an approach for deployment of e-services, we now describe several different technologies in various states of development to provide an overview of privacy-enhancing technologies for e-services. The reader may read further about broader guidelines for e-service or web service architectures elsewhere (WSA, 2004). The privacy technologies covered include: IBM's Enterprise Privacy Architecture (EPA), CONFAB: a system targeting pervasive services, a description of privacy features in the Liberty Alliance ID-Web Services Framework (Liberty Alliance-1, -2, -3) and an approach for using digital rights management to meet the needs expressed in the privacy principles. EPA is an architecture that is used by IBM in the development of privacy services for its clients. CONFAB is a research project that uses a policy-based approach for privacy management in pervasive applications. This work is of particular interest in the development of future e-services that would respect privacy preferences in an environment where a great deal of personal information, including location may be gathered and stored in many devices with different ownership. The description of the Liberty Alliance ID Web Services framework provides a real-world example of privacy management for a web service, whereas the digital rights management approach for privacy management describes procedures for handling several different situations dealing with personal data that can be applied for many different situations.

IBM Enterprise Privacy Architecture

The objectives for developing the IBM Enterprise Privacy Architecture were: helping organizations understand how privacy impacts business processes, and maximizing e-business trust. Based on privacy best practices and business requirements, EPA maps players, rules, and data to new or existing business processes by using object-oriented methods. This approach was intended to help organizations minimize the risks of inadvertent privacy disclosures by showing them where Personally Identifiable Information (PII) is stored in their enterprise and how to effectively manage it.

In order to introduce privacy-awareness and privacy services into enterprises in a systematic and complete way, IBM EPA is structured in four building blocks: the management infrastructure, the privacy agreement framework, the technical architecture, and the privacy regulation analysis. The management infrastructure enables an enterprise to define: its privacy strategy (e.g., embedding business best practices or rules into privacy policy), the general controls to enforce the privacy policy (e.g., supporting and ensuring general policy compliance), and the privacy practices to translate the privacy policy into its business processes. The privacy agreement framework provides a methodology for embedding the privacy policy into business processes, mapping the privacy parties, rules, and data, and minimizing the risks of inadvertent privacy disclosures. The technical architecture defines the supporting technology for implementing the required privacy services. The privacy regulation analysis identifies the applicable regulations. Figure 3 depicts the building blocks of the IBM EPA (Karjoth et al., 2002).

Figure 3. Building blocks of the IBM Enterprise Privacy Architecture.

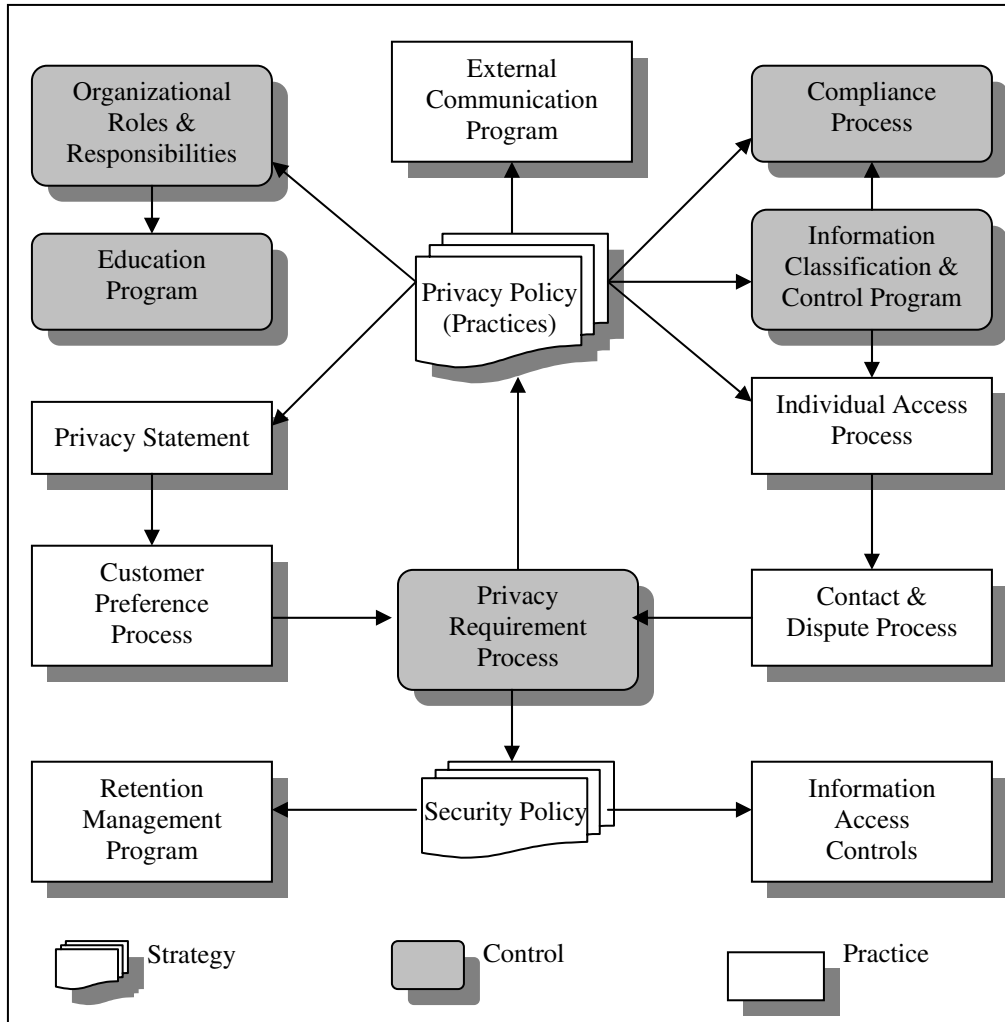


Management Infrastructure: The EPA management infrastructure is the tip of the EPA pyramid. It enforces an enterprise privacy strategy through a comprehensive privacy management program down to the implementation of privacy practices. The management infrastructure consists of three components: Strategy, Controls, and Practices. Figure 4 depicts their components and relationships (Brown, 2003).

- **Strategy:** This defines the high-level privacy and security policies and generates the privacy and security strategies for an enterprise.
- **Controls:** This defines the general controls to enforce the policies. The controls include a privacy requirements process, an information classification & control program, a compliance process, a definition of the organizational roles & responsibilities, and an employee education program.
- **Practices:** This incorporates and translates the policies into an enterprise's business processes. The practices include an external communication program, a privacy

statement, a customer preference process, an individual access process, a contact & dispute process, information access controls, and a data retention management program.

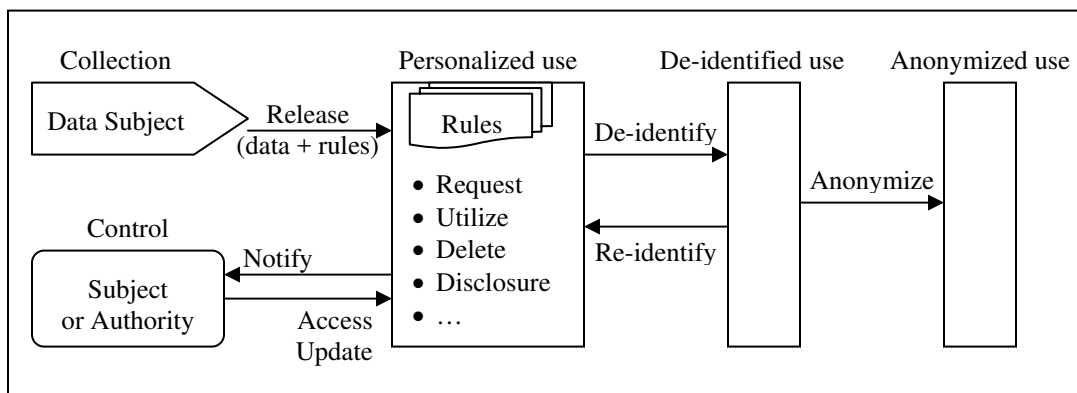
Figure 4. Components and their relationships of the EPA management infrastructure



Privacy Agreement Framework: The privacy agreement framework provides the privacy management for privacy-enabling business processes at the transaction level. The processes connect the individual to the enterprise, map data collection, storage, uses,

disclosures, and retention, minimize risk, and optimize the personal information (PI) handling processes by limiting collection, use, and disclosure according to the risk analysis of the threats and vulnerabilities. The framework consists of three major models: Players, Data, and Rules. Figure 5 depicts a process model for optimizing PI handling processes for privacy (Brown, 2003).

Figure 5. Process for optimizing PI handling processes for privacy



- **Players:** The players are the interaction entities in the data collection processing. They are data subjects or users.
- **Data:** This model identifies the required data for the collection processes. Based on the privacy-enhancing technologies (Goldberg et al., 2002; Lysyanskaya et al., 2000; Pfitzmann et al., 2000), the data can be categorized into three classes for privacy protection: personally identifiable information (PII), de-identified information, and anonymized information. PII is the most sensitive personal information that can be linked to a real-world identity, e.g., a social security number. De-identified information is the information replaced by a pseudonym. Anonymized information is

the least sensitive personal information that can be obtained by removing all personal data.

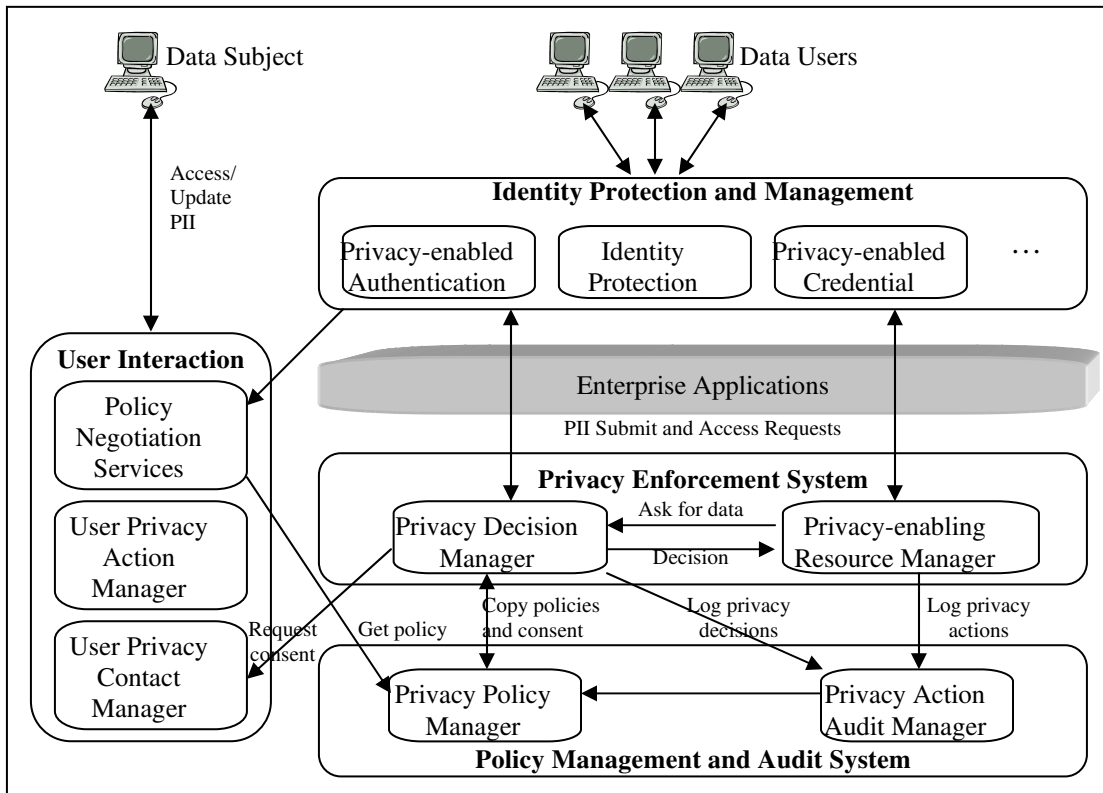
- **Rules:** This model identifies the rules for the data usage.

Technical Architecture: The technical architecture provides the necessary supporting technologies to ensure that an enterprise provides sufficient privacy protection to its customers. The technologies include a policy management system, a privacy enforcement system, and an audit console, and others. Figure 6 depicts the IBM EPA technical architecture (Brown, 2003). The major components are:

- **Policy Management System:** The policy management system enables the system administrators to define, change, and update policies, and assigns the policies to the privacy enforcement system.
- **Privacy Enforcement System:** The privacy enforcement system enforces the privacy protection for all personal data based on the policies it obtains from the policy management system, and offers the auditing data to the audit console. The privacy decision and privacy-enabling technologies are two major components of the privacy enforcement system. The privacy enabling technology can promise fair information practices to its customers. Furthermore, the privacy statements/policies can be formalized using the Platform for Privacy Preferences (P3P) and enforced directly within enterprise applications by the IBM Enterprise Privacy Authorization Language (EPAL).

- **Audit Console:** The audit console enables the privacy officers to review the policies and audit information.

Figure 6. IBM EPA Technical architecture



Privacy Regulation Analysis: EPA privacy regulation analysis uses regulatory summary and regulation rule tables to solve the regulatory compliance challenges. The regulatory summary tables summarize the applicable regulations with a unified terminology. The regulation rule tables describe the specific enterprise-based regulation rules with more formal style, for instance, an entry describing which party can access which type of data, and referring it to the legal regulation.

With privacy-friendly business processes and privacy-enabling security technology, IBM EPA provides a methodology for enterprises to enhance privacy protection for their customers. Its advantages include providing a well-defined level of privacy to customers, protecting the customer's data from privacy violations by regular employees, systems, or others, maximizing the business use of personal data for an enterprise, and respecting privacy regulations. In addition, the IBM EPA has integrated some new (e.g. EPAL) or existing (e.g. P3P) privacy-enhancing technologies into the system for defining enterprise privacy practices. IBM now uses EPA for its privacy technology consultation practice. It is also undergoing further research and development to create new privacy technologies related to: pseudonym-credential practice for identity protection, and privacy regulation analysis for privacy law/principles translation. The architecture assumes that the customers must trust the privacy administrators and privacy enforcement systems of the enterprise. Moreover, IBM EPA also appears to be the company's business model for targeting opportunities in their consulting services division for managing privacy for organizations (IBM Bus).

Privacy Architectures for Ubiquitous Applications and Privacy Policy

Compliance

In this section, we examine two previous works that deal with privacy architectures for e-services. Hong and Landay (2004) provide a toolkit, Confab, for facilitating the development of privacy sensitive ubiquitous computing applications. Yee and Korba (July 2004) propose an architecture for a privacy policy compliance system that operates within every service provider to ensure conformance to an user's privacy policy. In the

following, we summarize the key results of each work, and compare the approaches in terms of several headings.

CONFAB: Privacy for Ubiquitous Computing Applications

Hong and Landay address the difficulty of designing ubiquitous software applications that are privacy-sensitive or that help the user to manage his/her privacy (Hong & Landay, 2004). Their solution is to provide a toolkit with embedded data and programming models that can be used by developers to develop such applications. Summarized below are the privacy requirements for ubiquitous applications that they obtained through surveys of end users and application developers.

End-User Privacy Requirements:

- *Clear value proposition*: an upfront value proposition that leaves no doubt as to what benefits are offered and what personal information is needed to offer those benefits;
- *Simple and appropriate control and feedback*: simple control over and feedback about who can see what information about the end-user;
- *Plausible deniability*: addresses a social need to avoid potentially embarrassing situations, undesired intrusions, and unwanted social obligations, e.g. a person answers with a white lie when asked on the phone what they are doing;
- *Limited retention of data*: addresses concerns over long-term retention of personal data that can lead to unforeseen and unwanted use of the data;
- *Decentralized control*: addresses fear that personal data is stored on a central computer over which the end-user has very little practical control;

- *Special exceptions for emergencies*: the idea that in emergency or crisis situations, safety far outweighs privacy needs, e.g. disclosing personal health information in return for treatment in an emergency.

Application Developer Privacy Requirements

- *Support for optimistic, pessimistic, and mixed-initiative applications*: in pessimistic applications, end-users set up preferences beforehand and place strict constraints on when personal information can flow to others; optimistic applications allow greater access to personal data but make it easier to detect abuses after the fact with logs and notifications; in mixed-initiative applications, the end-user is interrupted when there is a request for personal information and he/she must make a decision to allow it or not on the spot;
- *Tagging of personal information*: marking personal information with privacy preferences, e.g. whether forwarding is allowed or amount of time to retain the information;
- *Mechanisms to control the access, flow, and retention of personal information (quantity)*: controlling the quantity of information disclosed to others, e.g. only people in the same building as myself can see my location, or colleagues can see my location between 9AM and 5 PM;
- *Mechanisms to control the precision of personal information disclosed (quality)*: granular control over the precision of disclosures (the quality of disclosures), e.g. giving one's location as "123 Main Street" or "Ottawa";

- *Logging*: for both clients and servers; for clients, logs facilitate understanding who is accessing what data; for servers, logs facilitate service audits to ensure that the clients' personal data is handled properly.

Hong and Landay summarize these requirements into four high-level requirements as follows:

- “A decentralized architecture, where as much personal information about an end-user is captured, stored, and processed on local devices owned by that end-user”;
- “A range of mechanisms for control and feedback by end-users over the access, flow, and retention of personal information, to support the development of pessimistic, optimistic, and mixed-initiative applications”;
- “A level of plausible deniability built-in”;
- “Special exceptions for emergencies”.


The authors illustrate the kind of applications that they wish to support with their toolkit by using two scenarios: a Find Friend scenario and a Mobile Tour Guide scenario. In the Find Friend scenario, employees can use a server to share their location information with one another. Employees choose to upload updates to the server at different levels, e.g. room level or floor level. The server is set up to notify a person whenever his/her location is queried and to accept queries only if the requestor is physically in the same building. In the Mobile Tour Guide scenario, a person visiting a city for the first time uses the Guide in conjunction with his/her location-enabled device to have a location-enhanced tour guide. The Guide can provide different levels of service depending on the level of

location detail the person shares. For example, if the location information is at the city level, the Guide can provide information on calendar events or the length of lines at major venues such as museums, whereas if the location information is at the neighborhood level, the Guide can additionally include information on interesting shops and other nearby points of interest.

To answer the above requirements, Hong and Landay devised the Confab toolkit with data and programming models that facilitate the design of privacy-sensitive ubiquitous applications. We next examine these models.

Confab’s data model makes use of “infospaces” that are assigned to people, places, things, and services in order to represent contextual information. An infospace is a network addressable logical storage unit that stores context data about the entity to which it is assigned (see Figure 7). For example, Alice’s infospace contains context information on her health, location, and activity. The following points apply to infospaces:

Figure 7. Ovals represent infospaces about a person, a place, or a thing. Squares represent tuples of contextual information associated with infospaces.



- Infospaces can be populated by sources of context data such as sensors.

- Applications can retrieve and manipulate infospace data to accomplish context-aware tasks.
- Individuals can specify privacy preferences for how their infospaces handle access control and flow.
- Infospaces are managed by infospace servers.
- The basic unit of storage in an infospace is the “context tuple”. Tuples can represent an attribute about an entity (e.g. a person’s age), a relationship between 2 entities (e.g. a person is in a room), static pieces of contextual information (e.g. an email address), or dynamic contextual information (e.g. a person’s location); tuples can optionally have a “privacy tag” that gives hints from the end-user on how that tuple should be used when it flows to another computer outside the end-user’s control (e.g. when the tuple should be deleted).

Confab’s programming model consists of methods and operators. Infospaces support 2 kinds of methods: “in” and “out”. In methods include add and remove, and determine the data stored in an infospace. Out methods affect the data leaving an infospace and include query, subscribe, unsubscribe, and notify. Infospaces also support operators for manipulating tuples. There are 3 types of operators: in, out, and on. In operators run on all tuples coming in through in methods, e.g. check the infospace’s access control polices to make sure the tuple can be added. Out operators run on all tuples going out through out methods, e.g. block tuples if end-user is in “invisible mode” (end-user does not want to give any information out). On operators run periodically, e.g. garbage collection. Hong and Landay include the following operator types (Table 2):

Table 2. Confab operators

Operator Type	Description
In	Enforce access policies Enforce privacy tags Notify on incoming data
Out	Enforce access policies Enforce privacy tags Notify on outgoing data Invisible mode Add privacy tag Interactive
On	Garbage collector Periodic report Coalesce

The interactive operator (Table 2) allows the end-user to have control over disclosures by displaying a simple GUI that allows the user to choose between disclosing the information just this once, ignoring it, or permanently denying access. The coalesce operator deletes tuples with repeated values. For example, location tuples can be duplicated if the person stays at a particular location over some period. The coalesce operator sorts the location tuples by time and deletes the tuples with duplicate values. Operators are loaded through a configuration file on startup and execute in the sequence in which they were added. Each operator in addition has a filter that checks if it should be run on a specific tuple. Once an in or out method is called, a sequence of the appropriate operators is put together and run on the set of incoming or outgoing tuples. Confab's programming model also supports service descriptions and active properties objects. Service descriptions are published by applications and provide basic information about the service as well as describe service options on features and what data types and data formats are needed from the user. A client application making a request on an infospace would first of all send its service description, for which the infospace can use a

previously stored configuration if it has seen the service before, or display a default GUI for the user to choose whether to allow access, choose options, and indicate how long the settings should last. An active properties object simplifies the task of querying for and maintaining context state in applications. Queries can be placed in an active properties instance and be periodically executed for up-to-date values.

We conclude our summary of Hong and Landay by describing one of the applications they built using Confab. They call this application “Lemming”, a new location-enhanced instant messenger client that provides two novel features. The first novel feature is the ability to request a user’s current location so that when the request is received, the end-user can choose “never allow”, “ignore for now”, “just this once”, or “allow if...” to allow requests under certain conditions. When a location request is received, the end-user’s instant messenger client issues a query to the user’s infospace for the user’s current location. The infospace checks to see if there is a context tuple representing location information, and then checks the age of the tuple to see if it is “current” (20 minutes by default). If the location tuple exists, it next flows through the out operators defined in the infospace. Three operators are of interest here: the Enforce Access Policies, the Interactive, and the MiniGIS operators. The enforcement operator checks if there is an existing policy associated with the end-user and applies the policy if it exists. The Interactive operator also checks for the policy and displays a GUI to let the end-user set a policy if the policy does not exist. The MiniGIS operator converts the data from latitude/longitude to a place name. The second novel feature involves the ability to automatically display a current location as an away message. The message can automatically update itself as the end-user’s location changes. The instant messenger

client sets up a query to retrieve the location every 60 seconds, and then displays this location in the away message.

The authors summarize the advantages of their work with the following points: an extendable suite of mechanisms for managing privacy, and personal information is captured, stored, processed on the user's computers.

Privacy Policy Compliance System

Yee and Korba examine how an e-service client can be assured that the e-service provider with whom he/she is interacting complies with his/her privacy policy (Yee & Korba, July, 2004). Underlying this is an e-services transaction model in which an e-service client and the corresponding provider each have a privacy policy that specifies their separate privacy preferences. The client's privacy policy specifies what private information the client is willing to give up and the conditions for access to the information (e.g. the provider can only have access during week days). The provider's privacy policy specifies what private information the service requires and the conditions that govern the provider's access to the information (e.g. need access every day of the week). The e-service can only be engaged if the client's privacy policy matches the provider's privacy policy. The authors' previous work considers policy negotiation (wherein there is no policy match) (Yee & Korba, Jan, May, 2003), how privacy policies can be semi-automatically generated (Yee & Korba, 2004), and how a match can be determined (Yee & Korba, 2005). However, we are interested here in their work on privacy policy compliance.

Yee and Korba's approach to the problem is to design a Privacy Policy Compliance System (PPCS) that has an embedded private data controller. Basically, the PPCS intercepts the user's data and ensures that processing of the data complies with the client's privacy policy. Prior to presenting the design of the PPCS, the authors derive requirements for the PPCS based on Canadian privacy legislation. These requirements are summarized as follows:

Requirements for the PPCS

- *Clear Purpose*: for each purpose for which private information is collected, the PPCS must provide clients with an explanation of what information is necessary in order to accomplish the purpose;
- *Limiting Use, Disclosure, and Retention*: for each purpose for which private information is collected, the PPCS must provide clients with an explanation of how it intends to use or disclose the client's private data; in addition, the PPCS must ensure that all copies (including copies disclosed to other parties) of the client's private information is deleted at the earliest of a) the time when the data is no longer needed for the fulfillment of the purpose, or b) the expiration of the data's retention time;
- *Accuracy*: the PPCS must provide a facility with which clients can access, check the accuracy, update, and add to their private data, as necessary for the corresponding purposes;
- *Openness*: upon request, the PPCS must display the provider's specific information about its policies and practices relating to the management of private information;

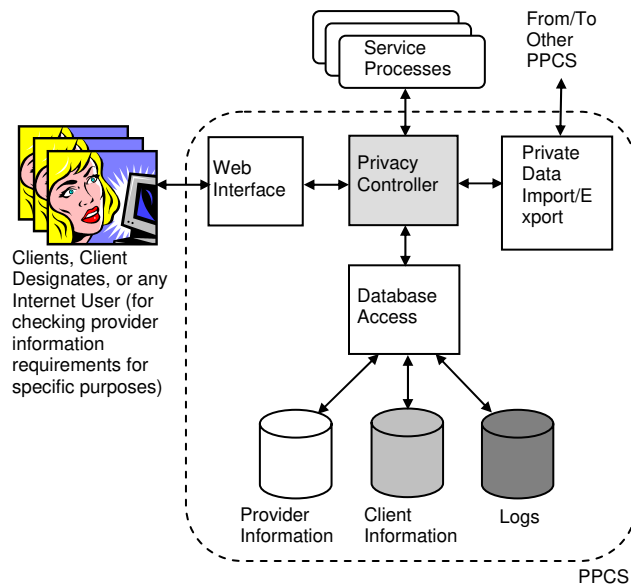
- *Individual Access*: upon a client's request, the PPCS must inform the client of the existence, use, and disclosure of his/her personal information, and give him/her access to that information;
- *Challenging Compliance*: upon request, the PPCS must allow the client or the client's designate to review the secure log (all PPCS actions are securely logged) to verify compliance to his/her privacy policy.
- *Safeguards*: the PPCS must have appropriate security safeguards in place to protect the client's private information from unwanted disclosure.

To satisfy these requirements, the authors propose the architecture in Figure 8.

We now describe each of the components in Figure 8. The *Web Interface* provides a UI for interactions with the client, client designate, or any Internet user (for checking provider information requirements for specific purposes). The *Web Interface* also establishes a secure channel to the client or client delegate and authenticates them. The *Privacy Controller* controls the flow of provider and client information and requests to fulfill the client's privacy policy; specific actions include: a) make log entries, b) delete private information upon completion of purpose or information expiry, c) grant access for client update of private information (including the update of information that has been provided to third party data processors), d) grant access for the examination of logs and comparisons of information, e) upon request, inform the client of the existence, use, and disclosure of his/her private information. The *Database Access* component provides read/write access to the databases as requested by the *Privacy Controller*, and handles

security protection for the databases. The *Private Data Import/Export* component sends private information disclosures to other providers, receives private information disclosures from other providers, sets up secure channels to other providers for sending information disclosures, and authenticates the providers. Three databases store information belonging to the provider, the client, and the system (logs). Provider information includes provider privacy policies, purposes, and so on. Client information includes client privacy policies and clients' personal information. Logs include entries for PPCS-client actions such as information collection, information use and disclosure, information access and update, and information deletion. Finally, the *Service Processes* represent the services offered; the arrow going out of these processes indicates private information collected by the services; the arrow going in represents private information required by the services.


Figure 8. PPCS architecture



Yee and Korba point out how parties who have received private information disclosures can be expected to delete the information upon completion of purpose or information expiry. “Such parties are considered to be subcontractor providers of the first provider and provide services to the first provider that are needed to complete the purposes of the first provider. In this case, the first provider is actually a consumer. As a consumer, the first provider has negotiated a consumer privacy policy with each subcontractor provider, containing the required purposes and information retention times reflecting the wishes of the original consumer. The PPCS of each subcontractor provider then deletes the original consumer’s private information upon completion of the purposes in the privacy policy agreed with the first provider or upon information expiry.”

The use of PPCSs to ensure privacy policy compliance is actually a distributed approach, where the PPCSs communicate among themselves to share information (via the Private Data Import/Export component). Each provider is expected to have one or more PPCSs, depending on how many services it is offering. This situation is depicted in Figure 9 (clients not shown).

Figure 9. Distributed nature of PPCSs. Here, Provider 3 discloses private information to Provider 1; Provider 5 discloses private data to Providers 2 and 4.



Comparison of Confab and PPCS approaches

We compare the above two approaches using the following 9 headings. Our comparison is given in Table 3.

- *Application area*: the type of e-services targeted by the approach;
- *Effectiveness at preserving privacy – general*: How effective is the approach at preserving user privacy? Is it foolproof?
- *Effectiveness at preserving privacy – disclosures*: How effective is the approach at preserving user privacy for user data that is disclosed to a third party? Is it foolproof?
- *Method for assuring clients*: How are users assured that their privacy has been preserved? What gives them the confidence?
- *Scalability*: Is the approach scalable?
- *System security*: Are any components of the implementation vulnerable to attack?

- *Validation*: Has the approach been validated through testing with a prototype?
- *Ease of implementation*: How easy is it to implement the system? Does it require further research? Are there already implementations or prototypes?
- *Costs*: How expensive is it to implement the system? Is the expense comparable to the expense of implementing similar software? Do the costs make business sense?

Table 3. Comparison of approaches for preserving user privacy

<i>Comparison Item</i>	<i>Hong and Landay</i>	<i>Yee and Korba</i>
Application Area	Ubiquitous application software, e.g. find friend service, mobile tour guide service; payment for service may not be first priority	Internet-based e-business, e.g. Amazon.com, Futureshop.ca; payment for service first priority; approach may be applied to distributed e-services
Effectiveness at Preserving Privacy - General	Fulfills privacy requirements under safe environment; can be defeated if infospaces and software not under user control	Fulfills privacy requirements under safe environment; can be made “mostly secure” (researching defense against malicious copying of user data)
Effectiveness at Preserving Privacy - Disclosures	Appears effective, making use of privacy tags, although details of security measures are not provided.	Appears effective, making use of recursive provider-client relationships, although details of security measures are not provided.
Method for Assuring Clients	Clients receive feedback, e.g. notification of location request; personally identifiable information stored on computing equipment owned by the client	Clients check secure logs to verify privacy policy compliance
Scalability	Mostly scalable, bottleneck may occur in high volume multiple disclosures to the same entity. There may be “policy chatter” caused by exchanges with entities requesting data.	Mostly scalable, bottleneck may occur in high volume multiple disclosures to the same entity or in exhaustion of a PPCS due to too many clients (but fix is to add more PPCSs)
System Security	Security measures not described but can be added	Security measures fully described
Validation	Validated by existing working applications	Needs to be validated by building and testing with a prototype
Ease of Implementation	Appears straightforward; privacy provisions are part of the design and present from the beginning	Appears straightforward but needs validation; PPCSs can be added to existing services
Costs	Incremental costs hidden in the costs of software development; does not appear to add inordinately more to the costs of	Very visible up-front costs (the cost of acquiring and adding a PPCS); however, fully recoverable as a cost of doing

	developing the same software but without privacy provisions	business due to attracting more clients through privacy provisions
--	---	--

Liberty ID-Web Services Framework – Privacy Features (Version 2)

The Liberty Alliance ID-Web Services Framework (ID-WSF) is an architectural platform for building secure, privacy-respecting, identity-centric web services. ID-WSF defines a common framework for web services of authentication, message protection, service discovery & addressing, policy & metadata advertisement, and data interaction (e.g. query & modify). More information is available from the Liberty Alliance (Liberty Alliance-1, -2, -3).

Privacy is a central tenet of the Liberty Alliance (there is an Expert Group within the Alliance dedicated to such issues) and ID-WSF in particular. The following sections highlight certain aspects of ID-WSF designed to enable good privacy. We provide a brief description below of some of the privacy considerations given for three different aspects of the service: consent, usage directives, and the interaction service.

Consent

ID-WSF-based entities may wish to claim whether they obtained the Principal's consent for carrying out any given operation. The Liberty SOAP Binding specification defines the **<Consent>** header block to allow Web Service Clients to indicate to the Web Service Provider that they have obtained the consent of the relevant Principal for the release of the location data.

The sample message below shows the <Consent> header block in a SOAP message requesting the release of a particular principal's location data.

```
<S:Envelope>
  <S:Header>
    <Consent id="A124395732495743"
      uri="urn:liberty:consent:obtained"
      timestamp="2112-03-15T11:12:10Z"/>
  </S:Header>
  <S:Body>
    Request for Location Data
  </S:Body>
</S:Envelope>
```

It is important to note that the **<Consent>** Header block as shown above is a claim made by the Web Service Client requesting the location data. The policy of the Web Service Provider hosting the service will determine if the claim is sufficient evidence of consent.

Usage Directives

The ID-WSF SOAP Binding provides a SOAP-based invocation framework for identity services. Within this framework, Liberty defines a usage directive container in which the policy requirements for attribute data, once released, can be carried. As an example, even if the privacy policy for a principal were to allow their PII to be released, the Web Service Provider might include with the location data any obligations that the requesting Web Service Client must fulfill or be in breach. Similarly, the Web Service Client can use the same **<UsageDirective>** Header block on its request to indicate its intent for the location data, if released. This is shown below.

```

<S:Envelope>
  <S:Header>
    <UsageDirective S:mustUnderstand="1">
      <cot:PrivacyPolicyReference>
        http://circle-of-trust.com/policies/eu-compliant/location
      </cot:PrivacyPolicyReference>
    </UsageDirective>
  </S:Header>
  <S:Body>
    Request for Location Data
  </S:Body>
</S:Envelope>

```

The Web Service Client inserts a reference to a specific privacy policy for location data in a PrivacyPolicyReference element (defined by some community of interest separate from Liberty). This information will feed into the Web Service Provider's decision to release the location data or not.

Interaction Service

A Web Service Provider will sometimes need to interact with the principal for which PII is being requested in order to clarify privacy policy. An Interaction Service allows Web Service Providers to ask the principal such as policy clarification queries without bearing the burden of maintaining the relevant addresses and details (e.g. Call me on my work phone during working hours but send me an IM Instant Message at any other times.)

Privacy Rights Management using Digital Rights Management

The examples above should manage the information flows appropriately. This section describes three typical message flows that must be maintained by privacy management architectures. Korba and Kenny described an architecture employing a rights management approach for the management of individual privacy rights as expressed by

European Union privacy principles. Their work goes on to provide some detail as to how to extend both XML (Kenny & Korba, 2002) and ODRL (Korba & Kenny, 2002) to meet the requirements for privacy rights management (PRM). This approach is useful in the context of systems like CONFAB and IBM's Enterprise Privacy Architecture as well as PPCS.

Within PRM there are four entities: the Data Subject (the person who owns the personal data), Personal Data (or Personally Identifiable Information (PII)), the Data Controller (the person, agency, public authority or other body which alone or jointly with others determines the purposes and means of processing personal data), and the Data Processor (the natural or legal person, agency, public authority or other body, which processes personal data on behalf of the controller). This arrangement logically matches the entities used to describe the obligations under privacy laws in many countries. Privacy principles are used to describe the general aspects associated with privacy laws (see chapter entitled "Legislative Bases for Personal Privacy Policy Specification" in this book). A privacy architecture must accommodate the privacy principles as they pertain to the service and jurisdiction in which it is being offered. In order to understand what is involved for Data Processors, Controllers and Data Subjects with respect to handling of personal data, one must explore the implications of the privacy principles and the systems involved. For instance, it is often the case that Data Subjects do not know which Controllers have what data, and whether it is accurate. Data Controllers and Processors may lose track of the data entrusted to them. This section explores the rights management approach by describing data flows related to particular data management cases and that are intended to meet privacy principle requirements.

Privacy Rights Management in Operation

Within PRM, servers handle the functions of the Data Controllers and Data Processors. In order to perform those functions, the Data Controller and Data Processor servers must maintain and use different sets of data. Below is a description of key controller and processor records and transaction logs required for PRM operation. These descriptions will facilitate understanding of the operational scenarios for data subject enrolment, periodic audit and personal data update by the data subject described in the following sections.

Processor/Controller Related Records:

Processors and controllers maintain 3 key record types regarding PRM operation. These include: processing agreements, audit information and Personally Identifiable Information Tracking data. Below, each is described separately.

- *Processing Agreements:* These are electronic documents containing the details of the arrangements between the controller and the processor. They contain information regarding: types of data the processor may accept, any limits to the processing prescribed by the Controller, time limits for access to data, agreements and details for audits (timing, type of data collected), as well as time stamp and approval signatures for the agreements.
- *Audit Information:* The Controller performs periodic audits of the data handling approaches for the processor. Results of the audits include a list of discrepancies between the data held by the processor as compared to those held by the controller. While detailed results are stored in the transaction log, the audit results for the

processor/controller are processed/summarized versions of those raw results for use by controller or processor.

- *PII Tracking Data:* The controller keeps track of the PII Data sent to each processor, the time of the transfer, and pointers to the policy and purpose for data processing.

Data Subject Related Records

There are several Data Subject-related records maintained by processors and controllers. These include the following.

- *PII Data:* The personal data entrusted to the controller by the data subject.
- *Contact Information:* Contact information for the PII Data (email address, home address, cross-referenced to PII Data, and policy and purpose for data use).
- *Audit Information:* Processed audit results pertaining to discrepancies in information regarding Data Subject data are stored here for review by the Data Subject.
- *Agreed-upon policies and purposes:* All privacy policies negotiated with the Controller and/or all Processors are stored along with a reference to the affected PII data.

Transaction Logs

In order to keep track of all activities of Data Controllers, Data Processors and Data Subjects within PRM, the following transaction logs are maintained.

- *Audit Results:* Detailed results from automated periodic, or external audits of the processes used by the processor and controller to assure PII is consistent and used only for the purposes and policies specified.

- *Transfers of PII:* Occurrences of transfers of PII. (timing, sender, receiver, and a reference to the PII involved).
- *Processing of PII:* All processors record time and duration of PII processing, as well as the policies exercised.
- *Policy Negotiation/Settlement:* Time of occurrence of privacy policy negotiation, with reference to the data subject, data processor, and/or data controller involved.
- *Data Subject Interactions:* Data Subjects may contact controller and processors to determine accuracy of PII data. Records are kept of all interactions.
- *Processor/Controller Interactions:* Timing and references to details pertaining to interactions between processors and controllers.

PRM Operational Scenarios

This subsection details PRM in operation by describing several key scenarios suggesting approaches within the PRM architecture intended to meet several key requirements of the privacy principles. The scenarios described here are:

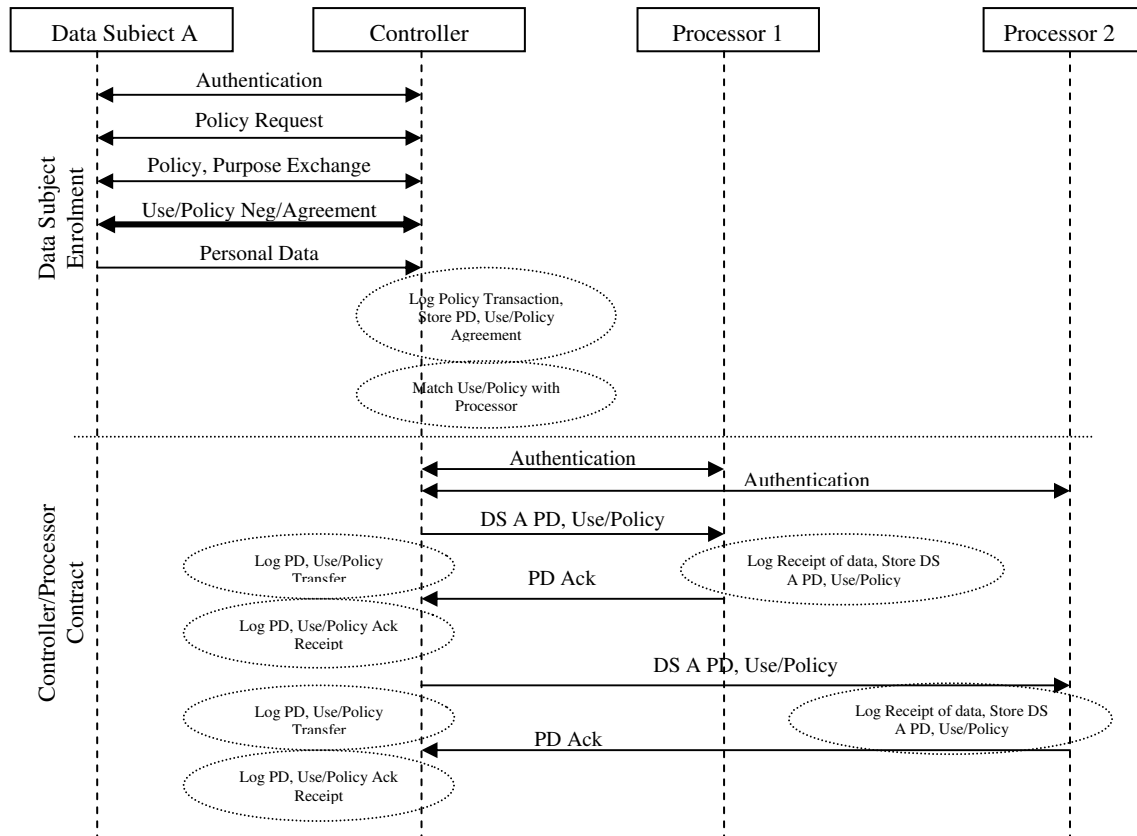
- Data subject enrollment,
- Periodic PII data audit, and
- Request for PII data update by the data subject.

The Scenarios are outlined using a description of data flows between the Data Subject, Data Controller and Data Processor within the PRM System.

Data Subject Enrollment

Data subject enrollment involves a data subject coming to agreement with a data controller on the personal information to be shared, as well as the privacy policy for dealing with the PII and the purpose for which the data may be used or processed. Figure 10 illustrates the data flow between the Data Subject, Data Controller, and two Data Processors.

Figure 10. Data flow during enrollment. (Personal Data (PD), Data Subject (DS), Acknowledgement (Ack))



The process starts with the Data Subject authenticating herself with the Data Controller.

For this and all further exchanges, the data subject and controller set up a secure

communication channel between themselves. The Data Controller exchanges a policy and purpose statement regarding the use of any personal data submitted by the Data Subject to the Controller. The Data Subject may negotiate with the Data Controller for a policy and purpose as described in (Korba, 2002; Yee & Korba, Jan., May, 2003). When the Data Subject comes to an agreement with the controller on the personal data to be exchanged, as well as the policy and purpose for which the data is being gathered, the data subject provides the data.

The Controller holds the personal data, exchanging it and the use and policy information with the processors that request the data. A number of log entries are made at various times during all of the exchanges. Figure 10 illustrates the various stages for enrollment in detail

Periodic PII Data Audit

Overseeing PII distributed amongst the Data Controller and Data Processors requires considerable effort and care on the part of the data controller. The Controller may have to deal with requests from data subjects or more detailed investigations conducted by a data protection authority. Either of these concern the quality of the data under the purvey of the Data Controller. Operating in a reactive mode to these investigations would be less desirable than a proactive approach wherein the Data Controller assesses the quality of PII under its purvey on a periodic sampled audit basis.

Figure 11. Data flow during Periodic PII Data Audit.

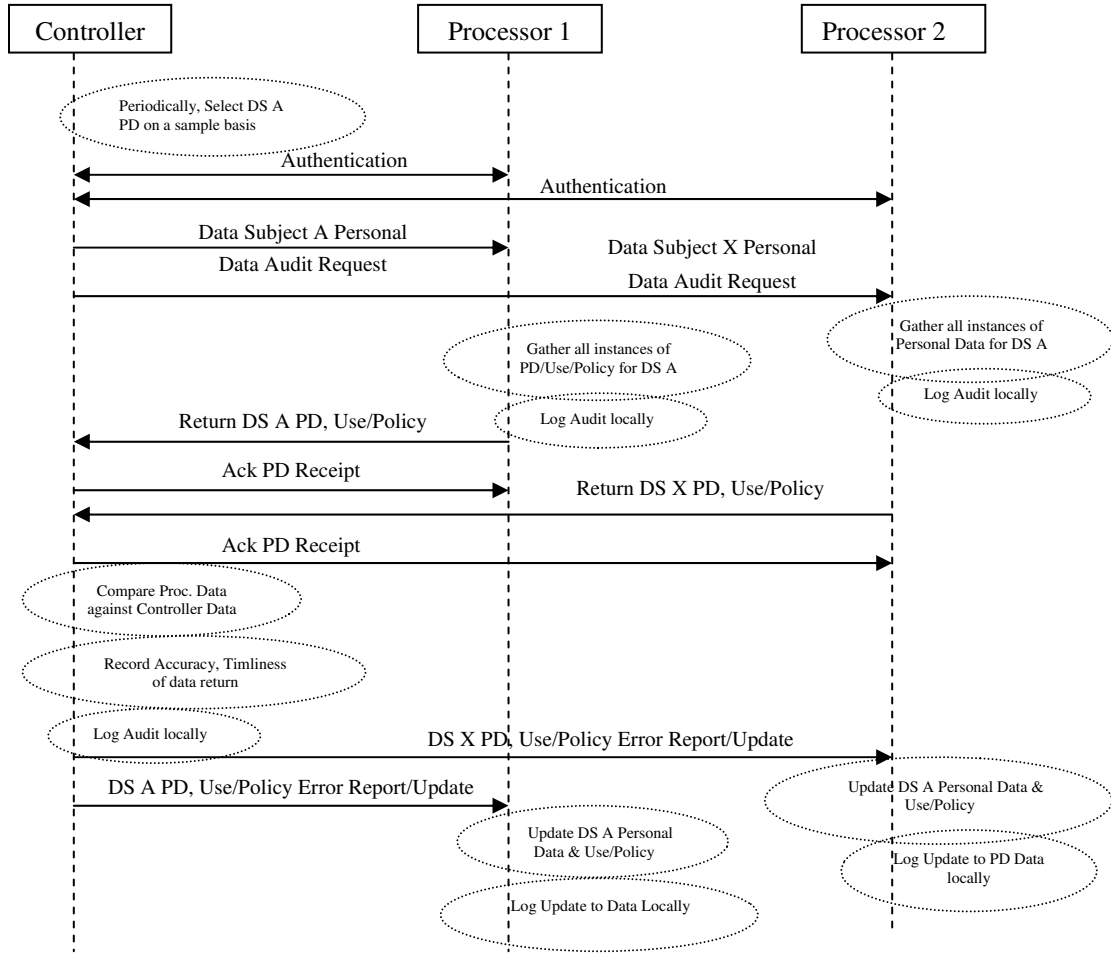
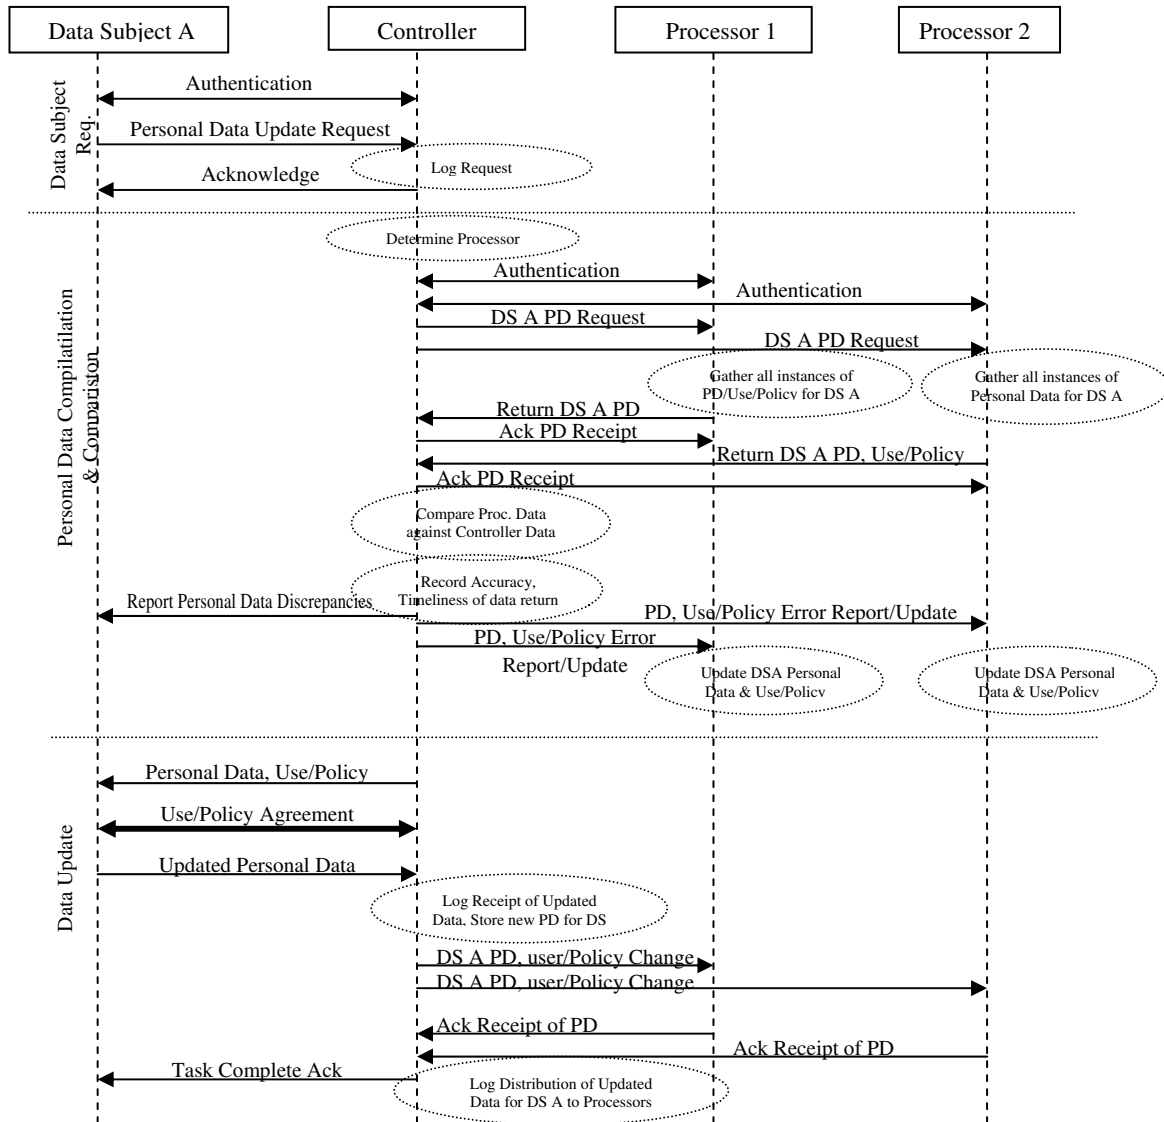


Figure 11 illustrates the interactions between the Controller and the Processors for the audit. The frequency of the audit would depend upon the amount of personal data held by the controller and processors and the desired level of quality. The Controller periodically selects personal data from different data subjects (shown as Data Subject X in Figure 11), polls all processors, requesting them to return personal data, policies, and purposes. The processors return the data (if any) they have for the selected Data Subject. From its records, the Controller determines whether or not the Processor should have the data, and

determines the accuracy of the personal data, policies, and purposes, by comparing them with its own records.

Figure 12. This Diagram shows the key interactions between Data Subject, Data Controller and two Data Processors during a user request for a change in Personal Data.



Personal Data Update by Data Subject

The Data Subject has the right to assurance that the data quality of their PII held by the Data Controller is maintained. The Controller may receive a request from a subject to check the data held by it. Figure 12 illustrates the update process. The Controller compares the data it distributed to the Processors against the original data received from the Data Subject in part to ensure there were no discrepancies in data handled by the different parties. Differences in Personal Data or Policies and Purpose are recorded and reported to the Data Subject. Any changes in PII requested by the Data Subject are made at the Data Controller and sent to the Data Processors that currently have the agreements with the Controller. The Data Subject may also negotiate policy and purpose for his/her PII.

FUTURE TRENDS

We started with a description of what is meant by e-services and a description of the driving factors behind privacy attitudes. We illustrated how privacy enhancing technologies are driven by corporate policies which are shaped by legislation and litigation, the development of new technology, consumer expectation, and competitive pressures. A privacy architecture may house many privacy enhancing technologies. Rather than prescribing a particular privacy architecture, we described several approaches that have been mentioned in scientific literature and presentations over the last few years. They included IBM's Enterprise Privacy Architecture, research projects Confab and PPCS, and the privacy architecture of the Liberty Alliance ID-Web Services Framework. We also provided examples of typical information flows that would be expected to be

supported in a privacy compliant architecture following on from research into a privacy rights management framework. All of these architectures are quite different from each other in implementation requirements. The Enterprise Privacy Architecture is an approach to deal with many aspects of IT operation, integrating some of the different tools IBM has developed (through its Tivoli arm), with legacy systems and a general approach for implementing and managing IT privacy. Confab is a research result targeting a privacy solution for the ubiquitous computing environment. It is especially relevant in the context of location sensitive services and e-services anywhere. The system was designed using the results of user and system developer surveys. PPCS is a research design that illustrates the approach of building a system from the ground up with legislated privacy requirements as the key driving forces in the design. The Liberty Alliance ID-Web Services Framework is an existing architectural platform that can be used today for building secure, privacy-respecting, identity-centric web services. The PRM work rounds out our examination of this area with examples of privacy-compliant data flows that would be supported by a privacy-enabled information system architecture. The common thread amongst these approaches is the use of a markup language to express rights and to track the use of data objects. Considerable effort is underway worldwide in the development of standards for different markup language variants to support processing of a wide variety of data in many different applications (e.g. Liberty Alliance). Many systems and inference engines have been developed and are currently under development for XML objects. In the future it may well be possible to use some type of XML technology to link between regulations, laws, privacy enhancing technologies, and privacy compliant architectures.

CONCLUSIONS

We have described some of the driving forces and approaches for the development and deployment of privacy architectures for e-services as well as presented several privacy information flow scenarios that can be applied for assessing privacy architectures.

Privacy management in e-services is a challenging multi-faceted task as demonstrated by the privacy management architectures we have presented. However, this challenge can be successfully handled using the architecture ideas and building blocks we have presented.

ACKNOWLEDGEMENTS

The authors would like to thank the National Research Council Canada for its support of this work. In addition, we would like to express our gratitude to Dr. Paul Madsen for contributing the material on privacy features of the Liberty Alliance ID-Web Services Framework.

REFERENCES

- Cranor, L. F., Reagle, J., & Ackerman, M. S. (1999). Beyond Concern: Understanding Net Users Attitudes about Online Privacy. *AT&T Labs-Research Technical Report TR99.4.3*, April 14, 1999, Available at <http://www.research.att.com/library/trs/TRs/99/99.4/>
- CSA Canadian Standards Association. Privacy Principles. Available at: <http://www.csa.ca/standards/privacy/code/Default.asp?language=English>

- Brown, N. (2003). Privacy Technology and the Public Sector. *12th CACR Information Security Workshop & 4th Annual Privacy and Security Workshop*, Toronto, Canada, November 6-7, 2003.
- Goldberg, I., Wagner, D., & Brewer, E. (1997). Privacy-Enhancing Technologies for the Internet. *IEEE COMPCON'97* (pp. 103-109).
- Hong, J. I. & Landay, J. A. (2004). An Architecture for Privacy-Sensitive Ubiquitous Computing. In *Proceedings of the Second International Conference on Mobile Systems, Applications, and Services (MobiSys2004)*, Boston, Massachusetts, June 6-9, 2004.
- IBM EPA. Enterprise privacy architecture. IBM privacy research institute. Available at <http://www.zurich.ibm.com/pri/projects/epa.html>.
- IBM Bus IBM Announces an Enterprise Privacy Architecture, June 29, 2001 Available at <http://www.bizwiz.com/bizwizwire/pressrelease/2005/8484ssw8x488ej7f88j.htm>
- Karjoth, G., Schunter, M., & Waidner, M. (2002). Privacy-enabled Services for Enterprises. In *Proceedings of the 13th International Workshop on Database and Expert Systems Applications (DEXA'02)*.
- Kenny, S. & Korba, L. (2002). Adapting Digital Rights Management to Privacy Rights Management. *Journal of Computers & Security*, Vol. 21, No. 7, November 2002, 648-664.
- Korba, L. (2002). Privacy in Distributed Electronic Commerce. In *Proceedings of the 35th Hawaii International Conference on System Science (HICSS)*, Hawaii, January 7-11, 2002.

- Korba, L. & Kenny, S. (2002). Towards Meeting the Privacy Challenge: Adapting DRM, *DRM 2002*, Washington, D.C., November, 2002.
- Lessig, L. (1998). The Architecture of Privacy. *In Proceedings of Taiwan NET'98*, Taipei, Taiwan. Available at http://www.lessig.org/content/articles/works/architecture_priv.pdf
- Liberty Alliance-1. Liberty Alliance Project. Available as of Feb. 28, 2005 from: www.projectliberty.org
- Liberty Alliance-2. Liberty Alliance ID-Web Services Framework Overview. Available as of Feb. 28, 2005 from: https://www.projectliberty.org/resources/whitepapers/Liberty_ID-WSF_Web_Services_Framework.pdf
- Liberty Alliance-3. Liberty Alliance & Privacy. Available as of Feb. 28, 2005 from: https://www.projectliberty.org/resources/trust_security.php
- Lysyanskaya, A., Rivest, R. L., Sahai, A., & Wolf, S. (2000). Pseudonym Systems. *In Howard Heys and Carlisle Adams (Ads.), SAC'99, LNCS 1758* (pp. 184-199).
- Pfitzmann, A. & Kohntopp, M. (2000). Anonymity, Unobservability, and Pseudonymity – A Proposal for Terminology. *In H.Federrath (ed.), LNCS Vol. 2009*, pages 1-9, Springer-Verlag, 2000.
- P3P. (2002). The Platform for Privacy Preferences 1.0 (P3P 1.0) Specification. *W3C Recommendation, April 16, 2002*. Available at <http://www.w3.org/TR/P3P/>.
- Solove, Daniel J., "Conceptualizing Privacy" . *California Law Review*, Vol. 90, p. 1087, 2002. <http://ssrn.com/abstract=313103>

WSA - Web Services Architecture Requirements, W3C Working Group Note 11

February 2004. Available at: <http://www.w3c.org/TR/wsa-reqs/>

Yee, G. & Korba, L. (July, 2004). Privacy Policy Compliance for Web Services. *In*

Proceedings of IEEE International Conference on Web Services (ICWS 2004),

San Diego, California, USA.

Yee, G. & Korba, L. (Jan., 2003). Bilateral E-services Negotiation Under Uncertainty. *In*

Proceedings of the 2003 International Symposium on Applications and the

Internet (SAINT2003), Orlando, Florida, USA.

Yee, G. & Korba, L. (May, 2003). The Negotiation of Privacy Policies in Distance

Education. *In Proceedings of 14th IRMA International Conference*,

Philadelphia, Pennsylvania, USA.

Yee, G. & Korba, L. (2004). Semi-Automated Derivation of Personal Privacy Policies. *In*

Proceedings of the IRMA International Conference, New Orleans, Louisiana,

USA, May 23-26.

Yee, G. & Korba, L. (2005). Comparing and Matching Privacy Policies Using

Community Consensus. *In Proceedings of the IRMA International Conference*,

San Diego, California, USA, May 15-18.

¹ NRC Paper Number: NRC 48271