Security Privacy Trust

International Security, Trust & Privacy Alliance

ISTPA

INTERNATIONAL SECURITY
TRUST & PRIVACY ALLIANCE

Privacy Framework

Version 1.1

Published by the International Security, Trust and Privacy Alliance.

The International Security, Trust and Privacy Alliance (ISTPA) has taken care in the preparation of the Privacy Framework™, but makes no expressed or implied warranty of any kind and assumes no responsibility for errors or omissions. No liability is assumed for incidental or consequential damages in connection with or arising out of the use of the information or programs contained herein.

For more information, please contact the ISTPA:

Address: 3525 Del Mar Heights Road, Suite 327
San Diego, CA 92130 USA

Phone: (858) 793-8100

Email: info@istpa.org

Website: http://www.istpa.org

ISBN: 0-9721484-1-8

Price $35.00 US

Table of Contents

Foreword

Digital technologies are rapidly and obviously transforming the way we conduct business and government, research and education, personal contacts and entertainment. The public is also becoming aware of the extent to which networked digital devices and databases are revolutionizing law enforcement, journalism, health care, litigation, and the activities of political parties, trade unions, charities, and advocacy groups – and also, unfortunately, of fraudsters, stalkers, thieves, and terrorists. Concerns about privacy go hand in hand with the expanded uses of personal information in public life.

Most of us enjoy the benefits of greater access to content and services, more far-reaching personal and professional contacts, more efficient, personalized, and convenient products and services, the promise of improvements in public safety. At the same time, we are uncomfortably aware of the increasing collection and diffusion of personal information that links us with a location, a history, a circle of family and friends, a level of income, a set of preferences or beliefs, a state of health or wealth. According to the opinion polls, most of us would like to be informed about how this information is used and exercise some control over at least the more sensitive uses. We would also like to be protected by law and technology against the more serious potential abuses, such as fraud and harassment.

As a consequence of these rising concerns, there has been a wave of new privacy-related legislation around the world, addressing specific dangers in some countries and in others establishing a broad regulatory scheme to protect privacy as a fundamental human right. There has also been a new emphasis in the marketplace on techniques and products for maintaining the security and confidentiality of data as it is collected, stored, and communicated, to meet the expectations of consumers whether or not they are reflected in legal obligations.

There is a fairly broad and international consensus on the principles, often termed “fair information practices,” that ethically (and often legally) must be taken into account in handling information about individuals. These include the notion that personal data should be collected and processed for defined and legitimate purposes, with fair notice to the individual, mechanisms for exercising choice (consent) wherever the individual’s privacy interests outweigh competing public or private interests, opportunities for access and correction by the individual, an appropriate level of information security, and practical methods of enforcement and recourse in the case of abuse or negligence.

But these principles of fair information practices are implemented differently, often with substantial variations in terminology, legal requirements, and technical and operational mechanisms, from country to country, sector to sector, and application to application. Privacy principles are a necessary foundation, but something more is required to achieve consistent, compliant, interoperable privacy and security solutions throughout an organization and across sectoral, state, or national borders. From a technology perspective, new tools have been developed which typically address very narrow privacy requirements, such as identifying instances of data collection through a website and providing notice of a website privacy policy, rather than assuring consistent and thoroughgoing privacy practices across an organization. Moreover, it is difficult even to frame the policy debate, at national level or within an individual company, over such specific and emotionally charged privacy issues as whether to require opt-in or opt-out forms of notice and consent and what limits should apply to various forms of covert information gathering and data mining, in the absence of a common framework for analyzing the privacy issues and options. In all the change, confusion, and emotion, a clearly defined and standardized set of operational privacy controls has not emerged. The ISTPA Privacy Framework is meant to fill that need.

The Framework is a product of extensive analysis of the fundamental constructs of information privacy – privacy principles and fair information practices, business and government data collection requirements, consumer and citizen rights, available technologies, and other relevant factors. Most of those who contributed to the Framework are engineers rather than policymakers – their chief objective was to express privacy functions in a practical and consistent manner so that they could be implemented in technology and operations, whatever decisions were reached as to the appropriate balancing of privacy and other interests in any particular case.

The Framework has been designed by a non-profit alliance of companies and organizations as a proactive tool which is able to support businesses in developing and managing their own privacy policies, even in the absence of law or regulation.

In version 1.1 of the Framework presented here, the fair information practices have been translated into a functional set of defined "services," some or all of which would be required to be performed in order to give effect to a privacy policy in any jurisdiction, for any particular application. Operational mechanisms, such as computer programs, can be written to perform the Framework services applicable to a given activity, and an organization's privacy controls can be measured for completeness against the Framework. Where a specific legal requirement applies, the appropriate Framework services can be employed to assure compliance. The Framework will presumably be of particular interest to organizations that must comply with privacy requirements in a variety of jurisdictions or applications, and to developers of privacy-oriented technology.

One of the tasks that ISTPA has set for itself now is to map some of the more important privacy legal regimes to the Framework, showing how it can be used for compliance purposes. Another is to show how the Framework services can be automated in some particular use cases. Our hope is that this will pave the way for practices that will work across a global company, for example, or the development of software solutions that can be applied to track and protect personal data across applications and users. Perhaps the Framework can serve as a sort of "middleware" bridging the differences in terminology and scope from one privacy regime to another and simplifying compliance with applicable laws, promises, and internal policies.

The Framework will not, of course, end the debate over appropriate policy choices that affect privacy. The privacy services defined in the Framework are themselves "policy-configurable" in each case. But the Framework may facilitate the identification of options and solutions in each case; it provides at least a common vocabulary and toolkit.

We welcome your review and comment on this work. We recognize it is just a beginning, and hope you will join with the ISTPA members to help refine and improve the Framework, making version 2.0 even more useful.

W. Scott Blackmer, Esq.

Bethesda, MD, October 2002

Member, ISTPA Board of Directors

John T. Sabo

Annapolis, MD, October 2002

President, International Security, Trust and Privacy Alliance

Introduction

Overview

For many years, security practitioners have benefited from the existence of specific technology frameworks and a recognized, even standardized, set of security services and security technologies. These include cryptography, defense-in-depth architectures, public-key infrastructure, and secure communication protocols – all developed to address specific security requirements and to mitigate security risk.

Similarly, pressures now exist to provide an enabling, technical infrastructure for the provision of *privacy* in a business context as multiple forces interact to magnify concerns for privacy in modern society. The speed, ease of use, and ubiquitous nature of the Internet have made the gathering and distribution of personal information almost instantaneous. The growth of centralized and distributed computer systems, databases, and data mining technologies combined with the development of extensive networks for information exchange, have added new dimensions to the challenges of managing data privacy.

When sophisticated technical capabilities are combined with competitive business pressures to "know and capture the customer," the capability to meet consumer, citizen and legislated privacy preferences and mandates becomes increasingly difficult, and the potential for misuse and abuse of personal information and the subsequent loss of trust become major worldwide public policy issues.

The purpose of this document is to present a Privacy Framework that will provide a way to combine privacy and security throughout the life cycle of personal information.

ISTPA understands that the collection and processing of personal information are essential to the proper functioning of modern society and commerce. ISTPA believes that a high-level framework can be used as a base from which supporting architectures, technologies, tools, and complementary operational practices can be developed to support any set of privacy requirements, including cross jurisdictional policies. Such a framework consists of operational yet optional modules that are configured with the specific privacy policies and relevant parameters in each particular context. Contextual data is not "hard coded", but rather the framework is *policy-configurable*. A framework will enable businesses to deploy, for the first time, automated mechanisms which will support the core definition of information privacy: the proper handling and use of personal information throughout its life cycle, consistent with data protection principles and the preferences of the subject.

This is groundbreaking work, and ISTPA is pleased to contribute the Framework as a catalyst for what we believe will be a growing, formal body of technical work to advance the state of information privacy.

Executive summary

If information privacy is the proper handling and use of personal information throughout its life cycle, consistent with data protection principles and the preferences of the subject, then p*ersonal information* (PI) is any data related to an individual or entity, regardless of whether the subject of the PI is identified. Worldwide, especially with the rapid onset of web-based e-business, privacy concerns have intensified and legislation has been enacted that mandates stringent behavior in dealing with PI. A policy-configurable framework will allow the particular jurisdictional requirements to be input as parameters that then govern the behavior of the framework.

For more than 30 years, a set of principles and *fair information practices* have been evolving in the business and government sectors for the handling of personal information. These practices include:

- Notice and awareness
- Choice and consent
- Access (by the subject of the personal information)
- Information quality and integrity
- Update and correction, and
- Enforcement and recourse.

These practices serve as high-level guidelines for human and computer system behavior toward PI, but the operational specifics are left to the implementer.

The ISTPA Privacy Framework consists of seven services and three capabilities that faithfully implement the fair information practices, but which contain operational details. The seven services are **Audit, Certification, Control, Enforcement, Interaction, Negotiation, and Validation;** the three capabilities are **Access, Agent and Usage.** A capability is a virtual service that derives its functionality by “calling” other services.

Use cases illustrate how the various mechanisms within each service or capability can be exploited in specific contexts.

The Framework can serve as a template for designing privacy management systems and as an analytic tool for assessing privacy solutions. The Framework services and capabilities can be combined with existing, industry-standard security architectures to create a robust information privacy solution that can be tailored within and across jurisdictions.

Audience

This document is primarily intended for the privacy officer, or those persons responsible for implementing privacy policies and controls in organizations. At the same time, this document can be used to stimulate additional and more technical specification work, especially in appropriate standards bodies. Since the document is largely self-contained, a general audience of information technology and business professionals would also benefit. Legislators and government officials can reference the ISTPA Privacy Framework as they work with legal issues to determine how any new or revised legislation in this arena complements what is technically feasible.

Organization of this Document

The ISTPA Privacy Framework is composed of a set of ten privacy services and capabilities expected of any system for privacy management. After a brief introduction to the motivation, each service or capability is defined in turn, along with use cases that demonstrate how the service or capability interacts with the rest of the Framework.

Why use this document

This document specifies a Privacy Framework that can be used to assist in supporting privacy principles, implementing the privacy fair information practices, and developing architectures and technical implementations needed to support business privacy policies and consumer agreements. Operational privacy services and capabilities are outlined that can guide programmatic and policy development and serve many other related purposes.

The Framework will evolve toward more detailed functional descriptions of the services and capabilities and their interactions. Additional use cases, some with a specific industry orientation (e.g., financial services, CRM, eGovernment, location-based services, wireless medicine, mobile commerce and Enhanced 911), will be developed to further test and possibly modify the Framework.

The ISTPA welcomes additional companies to join our expert working groups and to advance the vital follow-on work suggested by the Framework document. Membership information can be found at www.istpa.org, the ISTPA website.

In order to make the benefits of the Framework available generally to industry and government, ISTPA publicly releases this document and encourages feedback and comments. Send your comments to director@istpa.org.

About the ISTPA

The International Security, Trust, and Privacy Alliance <www.istpa.org> is a global alliance of businesses and technology providers. Our goal is to work together to provide objective and unbiased research and evaluation of privacy standards, tools, and technologies, and to define a privacy framework for building technology solutions.

Within the ISTPA, there are several working groups, each with a particular area of focus. The Framework Working Group is responsible for developing and promoting an objective framework for achieving security, privacy, integrity, and trust in all forms of communications worldwide, as described in this document.

The ISTPA Privacy Framework can be used as a guideline or template for developing operational solutions to privacy issues and as an analytical tool for assessing the completeness of proposed solutions. The Framework is not yet a “specification” in the formal sense, but can be evolved into a specification.

Acknowledgments

The ISTPA Privacy Framework is a joint-volunteer effort made by many ISTPA members who provided insight, commentary and direction. A special thanks and recognition is due to the Framework Working Group and the contributing authors who patiently and diligently created and shaped the content and who collaborated to articulate and design the purpose, benefits and vision of the ISTPA Privacy Framework. Critical to the Privacy Framework is a multi-disciplinary and unifying approach, recognizing that the problem domain - the challenges of security, trust, and privacy - requires a multi-disciplined approach. If the Privacy Framework is to be successful, it must remain a collaborative and global effort built with careful attention to the diverse issues and complex technologies that our global information society and digital economy struggle to integrate and resolve.

Our gratitude goes out to these contributing authors and to the many future contributors to the ISTPA Privacy Framework:

Framework Working Group

Michael Willett, Chair, Wave Systems
Professor Thibadeau, Co-chair, Carnegie Mellon University
Lark Allen, Wave Systems
Kevin O'Neil, CYVA Research
Drummond Reed, OneName
Gary Roboff, GSR Strategic Consulting
Geoffrey Strongin, AMD
Betty Whitaker, NCR (formerly)
Monty Wiseman, Intel

Framework Editorial Team

Anne Jackson, Managing Editor
Michael Willett, Wave Systems
Kevin O'Neil, CYVA Research
Monty Wiseman, Intel

ISTPA Board of Directors

The ISTPA Board of Directors volunteers its time to advance the organization's projects and support the several ISTPA Working Groups. ISTPA Board members and officers, when Version 1 of the Framework was finalized, included:

Officers

John Sabo – President, Computer Associates International
Gary Roboff - Vice President, GSR Strategic Consulting
John Lindquist – Treasurer, EWA IIT
Cheryl Charles – Secretary, BITS

Additional Board Members

Scott Blackmer, Privacy Attorney
Michele Drgon, Motorola
Jim Schreckengast, Gemplus
Drummond Reed, OneName
Geoffrey Strongin, AMD
Norrie Taylor, NCR
Michael Willett, Wave Systems

Executive Director

Kevin O'Neil - CYVA Research

Special thanks go to Ronn Bailey, Founder of Vanguard Integrity Professionals, Inc. and former Board member, who generously supported the ISTPA website and email servers while the ISTPA Privacy Framework was under development.

Memoriam

In memoriam, we dedicate this work to George Jelen. George was especially devoted to the work of the Framework Working Group, contributing content and insightful commentary. He kept us focused.

High-level overview of the Privacy Framework

A potential solution to privacy management would be a collection of behaviors that faithfully satisfy the mandates in the definition of privacy, within a wide variety of contexts and scenarios. Privacy is the proper handling and use of personal information (PI) throughout its life cycle, consistent with data protection principles and the preferences of the subject. Since PI has a life cycle, the implication of the definition is that *proper* and *consistent* apply throughout the PI's life cycle.

The fair information practices constitute a collection of behaviors, but do not by themselves contain a structural framework for defining specific and repeatable functions.

From a high-level, operational viewpoint, start with the essential elements of the definition of privacy: proper handling, use, consistency, and preferences, with the focus on privacy management throughout the life cycle of personal information.

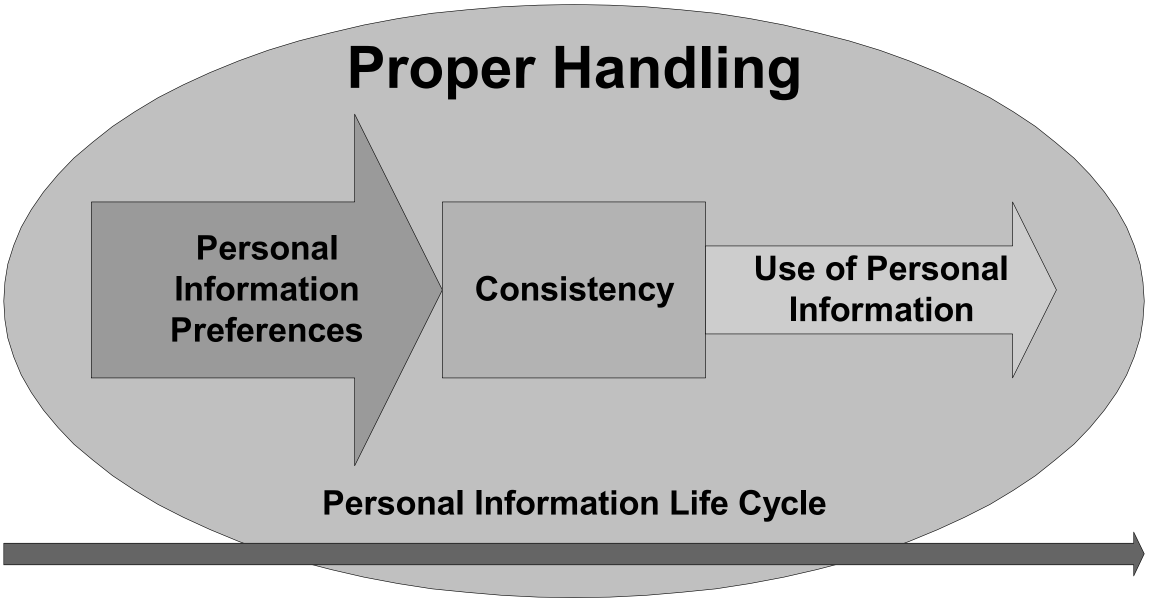
Proper Handling
Personal Information Preferences
Consistency
Use of Personal Information
Personal Information Life Cycle

Privacy Management

The ISTPA Privacy Framework provides a solution that maps the fair information practices to the challenges shown in the figure above. Privacy policy and other operational parameters are input to the Framework on a contextual basis in support of policy-configurable and adaptable implementations.

Privacy principles and fair information practices

This section introduces the privacy principles and fair information practices that serve as the design requirements for an operational privacy framework. The definition of privacy is reflected in the practices across the life cycle of personal information to produce a set of privacy services and capabilities. A simple use case will illustrate the operational aspects of the ISTPA Privacy Framework.

Background

The business scope of privacy concerns applies to transactions both on and off the Web, including traditional "brick and mortar," and newer "click and mortar" and e-commerce and e-government environments. The ISTPA Privacy Framework must apply to all business scenarios that involve the potentially improper collection and use of personal information.

Security and privacy concerns continue to be the leading inhibitors to consumer participation in e-commerce. Consumer concern is heightened by high-profile publicity involving mishandling of personal information and failure of security safeguards. In the United States, there are continuing pressures for legislative action, even as the business community is encouraging self-regulation. In other countries, there is a growing sense that implementation of legislated privacy rules has fallen behind the technology and that policy enforcement is difficult both for business owners and privacy officials. Ultimately, building trust between consumers and businesses depends on delivering value for both, while providing sustainable security and properly safeguarding personal information throughout its life cycle. Corporate policies and procedures in security and privacy require careful design and systematic deployment, supported by corporate management. Considerable risk reduction and brand protection can result from such deployment. **The bottom line is trust**.

The ISTPA Privacy Framework provides a technical, administrative, and legal template for developing a trustworthy infrastructure satisfactory for business, government and consumers.

Fundamental definitions

There have been many different viewpoints about privacy. Over a century ago, United States Supreme Court Justice Louis Brandeis defined privacy as "the right to be let alone," which he said was one of the rights most cherished by Americans. With work undertaken in the United States in the 1970s and enhanced in Europe in the decade of the 1980s, information privacy emerged as an area of study and definition.

Today we often see confusion between security and privacy. This confusion has served to constrain the development of trusted infrastructures that address both disciplines. For purposes of developing the Framework, ISTPA adopted these business-based working definitions of privacy, security and their relationship:

- **Security** - the establishment and maintenance of measures to protect a system.
- **Privacy** - the proper handling and use of personal information throughout its life cycle, consistent with data protection principles and the preferences of the subject.

- **Personal information (PI)** - information related to an individual or to entities other than individuals (for example, corporate entities may have a concern about "personal" information related to the corporation).

Security is necessary for privacy, but the proper handling and use of personal information requires an even broader set of privacy management functions.

Privacy principles

Privacy principles were first articulated in a comprehensive manner in the United States Department of Health, Education and Welfare's 1973 report entitled "Records, Computers and the Rights of Citizens." In the years since, sets of privacy principles have been developed by a variety of governmental and inter-governmental organizations. Key reports describing the core principles include "The Privacy Protection Study Commission, Personal Privacy in an Information Society (1977)," and the Organization for Economic Cooperation and Development's "OECD Guidelines on the Protection of Privacy and Transborder Flows of Personal Data (1980)."

Generally accepted privacy principles include:

- **Accountability** - ability to address the improper handling of personal information.
- **Collection limitations** - limiting the types of information collected.
- **Disclosure** - informing the subject when personal information is collected.
- **Participation** - allowing subject choice over collection and distribution.
- **Relevance** - collecting only personal information pertinent to the application.
- **Security** - protecting personal information from unauthorized access, alteration or destruction.
- **Use limitations** - limiting the subsequent use of collected information.
- **Verification** - checking the validity of personal information.

The privacy principles are high-level design points that describe the proper handling of personal information.

Fair information practices

At a more "operational" level, Fair Information Practices represent definable actions that are necessary to support privacy principles. A set of practices that are now widely accepted include (U.S. Federal Trade Commission):

- Notice and awareness
- Choice and consent
- Access (by the subject of the personal information)
- Information quality and integrity
- Update and correction
- Enforcement and recourse

The fair information practices can be used as a guide in implementing the privacy principles, since the practices are more operational in nature. However, even the practices are missing essential elements to support a technical, programmatic implementation such as subject agent, interfaces, policy control, and secure repository. The ISTPA Privacy Framework provides a more complete template for an implementation of the practices by including the missing elements. At a more granular level, the Framework allows specific technologies for the various functions to be selected, as appropriate for the given environment.

In summary, the ISTPA Privacy Framework accurately translates both the privacy principles and practices into a lower-level embodiment of privacy services and capabilities in such a way that specific implementation mechanisms are suggested.

Life cycle management of personal information

The life cycle of personal information assumes a period of time when the data subject may not have immediate or physical control over the information. Yet, the subject may desire to maintain strong, vicarious control over subsequent transfer and usage, or may wish to review, modify or withdraw any agreements. The agreement, perhaps including transfer and usage rules or limits on re-disclosure, should be robustly linked or associated with the personal information, so that subsequent processing actions will be managed in accordance with the agreement of the parties.

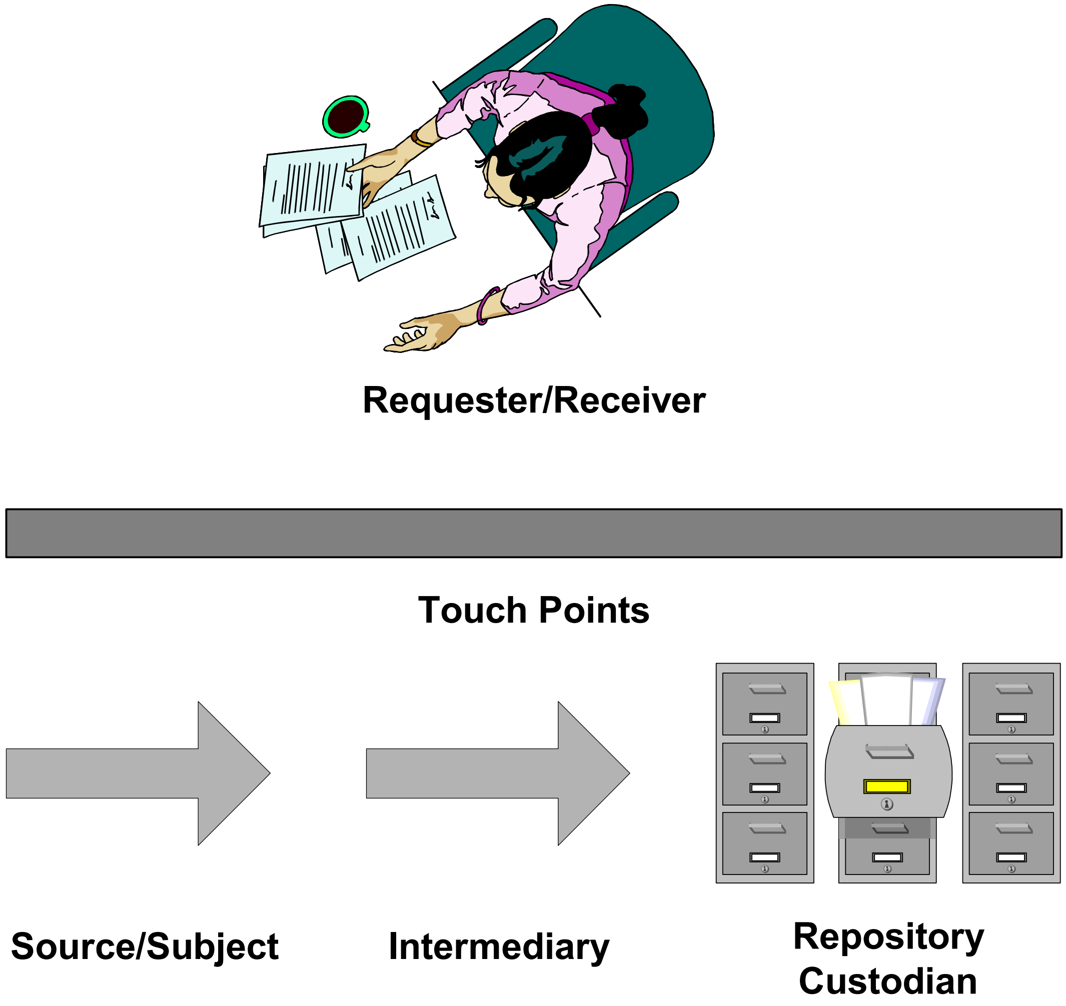

Life cycle of PI Management

PI touch point structure

To describe the relationship between the PI and the policies and agreements governing its use, a touch point can be structured, with the personal information at the innermost protected level. The integrated security and privacy service layer provides the necessary protection and proper handling of information, supporting in detail the privacy principles and practices at a functional level. This technology layer is parameterized, so that specific technology, policy, and jurisdictional choices are not pre-selected. Instead, the next layer is a legal, technical, and administrative layer that allows the touch point to be configured for various technologies, jurisdictions and organizational constraints. Since the Framework is policy-configurable, not all implementations will support the same functions or exhibit the same behaviors, since the parameterization layer supports jurisdictional customization. Finally, the requestor or receiver of personal information interfaces to these underlying layers to guarantee appropriate and accountable use.

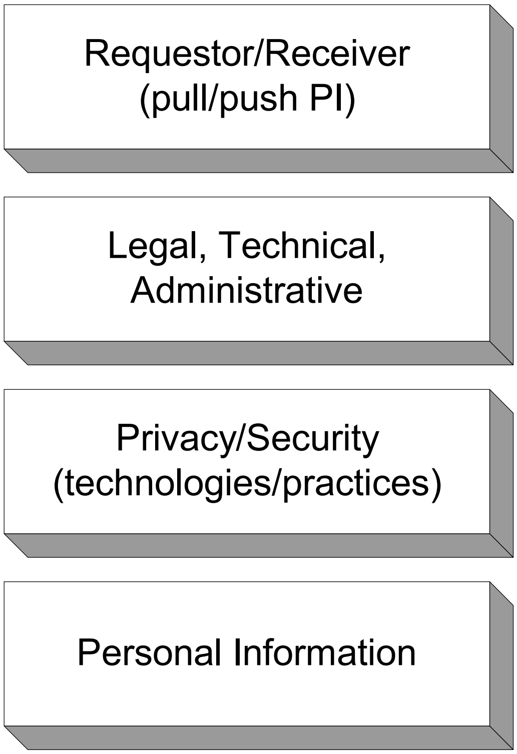
Requestor/Receiver (pull/push PI)
Legal, Technical, Administrative
Privacy/Security (technologies/practices)
Personal Information

PI Touch Point Structure

Security requirements and services

Although the Framework does not directly address security requirements, security is an essential component of information privacy. Security services are mandatory throughout the lifecycle of PI. The privacy services and capabilities developed in this document can be complemented with any standard set of security services to satisfy the privacy requirement for security.

For example, The Open Group extended the International Standards Organization (ISO) architecture (ISO 7498-2 Security Architecture) for security, expanding the earlier, “protectionist” view of security (security as defense) in order to develop a structure for enabling security techniques in support of e-business requirements (security for authorized access, for digital signatures, etc.). The result is an international standard for a security architectural framework, called the Architecture for a Public Key Infrastructure (APKI: pub 801). In this model, security requirements were taken through a layered analytic approach to deduce the needed security services and underlying mechanisms. A similar approach was used by the Framework Working Group to develop the ISTPA Privacy Framework.

The Open Group's business approach to security is complementary to the ISTPA Privacy Framework and in fact can serve as a model for mutually exploiting privacy and security services using the framework model. However, other security models can also satisfy the security requirement for privacy.

Privacy services and capabilities

The exercise conducted by the ISTPA Framework Working Group has evolved the following list of privacy services and capabilities, based on the requirement to support the privacy principles and practices described above, but at a functional level. A system and process design person should be able to integrate these privacy services/capabilities into a functional architecture, with specific mechanisms selected to implement these functions. In fact, the purpose of the ISTPA Privacy Framework is to stimulate design and analysis of the specific functions, both manual and automatic, that are needed to implement the complete set of privacy fair information practices. In that sense, the ISTPA Privacy Framework is an analytic framework.

To create a usable framework, various system capabilities must be identified that are not explicit at the privacy practices and principles level. For example, a policy management (or control) function is essential to honor the PI usage constraints established by the subject, but such a function is not explicitly called out in the privacy principles. Likewise, interfacing to the Framework is not explicit in the privacy principles, but is another essential operational service. Such inferred services are necessary if information systems are to be made “privacy aware.” Without them, enforcing privacy requirements in a fully automated environment will not be possible, and both businesses and consumers will be burdened with inefficient and error-prone manual processing.

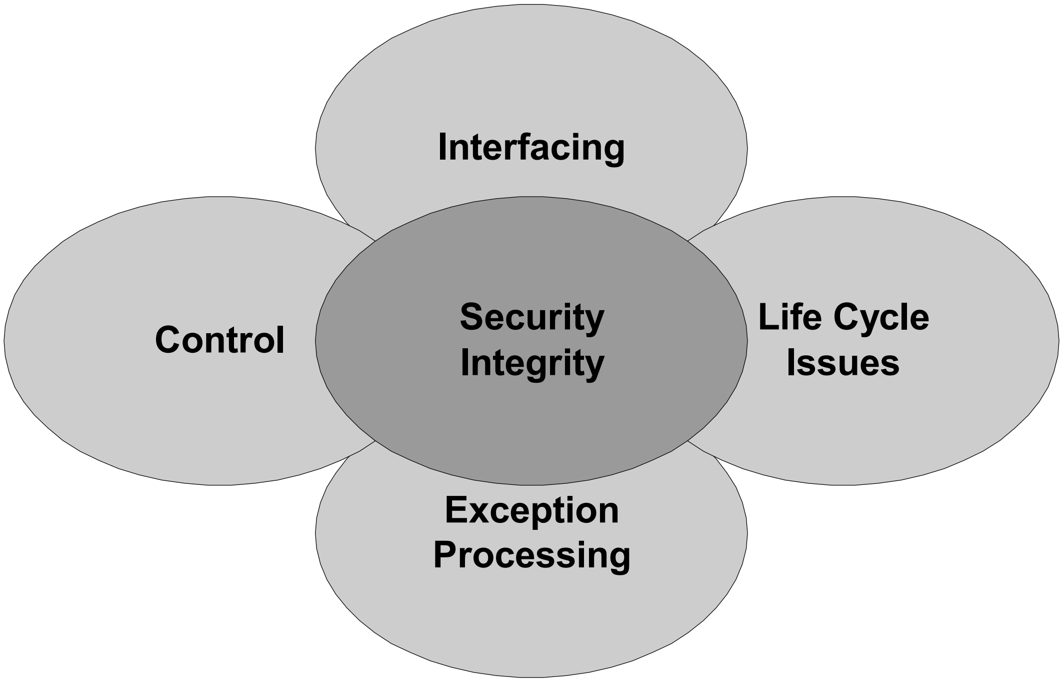

Operational Requirements

The ISTPA Privacy Framework identifies seven privacy services and three capabilities:

- Audit
- Certification
- Control
- Enforcement
- Interaction
- Negotiation
- Validation
- Access (capability)
- Agent (capability)
- Usage (capability)

A “service” is a collection of related functions and mechanisms that operate for a specified purpose; a “capability” also operates for a specific purpose, but does so by invoking multiple services. Abstractly, services and capabilities operate at the same logical level in the architectural hierarchy.

Summary of services and capabilities

The following chart summarizes the services and capabilities in the ISTPA Privacy Framework. The descriptions have been derived by examining the privacy principles and fair information practices and by identifying the necessary operational functionalities.

The totality of functions needed to realize the fair information practices was divided in a natural way to create the services and capabilities. This division into functional groupings is not unique, but the robustness and usefulness of this particular choice will be tested with use cases and further business and technical community analyses.

Service / Capability	Description
Audit	Handles the recording and maintenance of events in any service to capture the data that is necessary to ensure compliance with the terms and policies of an agreement and any applicable regulations.
Certification	Manages and validates the credentials of any party or process involved in processing of a PI transaction.
Control	Functions as “repository gatekeeper” to ensure that access to PI which is stored by a data collection entity complies with the terms and policies of an agreement and any applicable regulations.
Enforcement	Handles redress when a data collection entity is not in conformance with the terms and policies of an agreement and any applicable regulations.
Interaction	Presents proposed agreements from a data collection entity to the data subject; receives the subject’s personal information, preferences, and actions; confirms actions; manages movement of data into and out of the Framework. To the extent the data subject is represented by an

	agent, this service comprises the interface to the agent.
Negotiation	Handles arbitration of a proposal between a data collection entity and a data subject. Successful negotiation results in an agreement. Humans, agents, or any combination, can handle negotiation.
Validation	Checks for accuracy of PI at any point in its life cycle.
Access	A capability that allows the data subject to both access the individual's PI that is held by a data collection entity, and to correct or update it as necessary.
Agent	A software capability that acts on behalf of a data subject or a requestor. The Agent Capability engages with one or more of the other services defined in this Framework. Agent can also refer to the human data subject in the case of a manual process.
Usage	Functions as "processing monitor" to ensure that active use of PI complies with the terms and policies of an agreement and any applicable regulations. Such uses may include transfer, derivation, aggregation, pseudo-anonymization, linking, and inference of data.

Specific operational behavior of these services/capabilities is governed by the privacy policy and parameters configured in each jurisdictional context.

Binding Service / Capability Interactions with The PI

In order to transport the PI that is bound or associated throughout its life cycle to the agreed policies and permissions, a "PI container" is described. Such a container object and the binding mechanism are abstract constructs for purposes of the Framework. In effect, the binding mechanism is a configuration parameter, from simple pointers to full cryptographic binding. Included in the container are the subject/receiver contract, including any negotiated permissions, and the credentials for the subject, plus other contextual information that applied at the time the association was made. The chart below shows a representative PI container and the binding element.

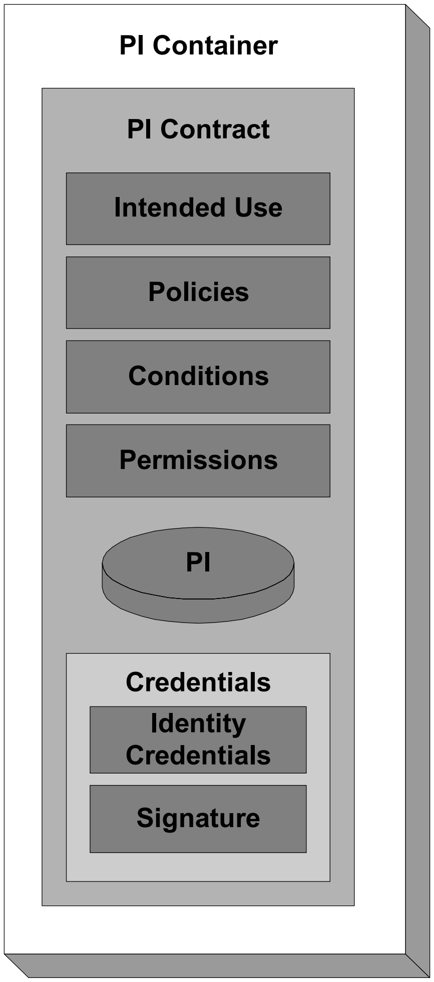

PI Container

Shown in the next chart is a logical configuration of the services and capabilities in the ISTPA Privacy Framework, with an Agent Capability representing both the subject and the data requestor. Interaction, Negotiation, and the Control Service provide a front-end to the secure data repository. The assurance services of Validation, Certification, Audit, and Enforcement support both subject and requestor, whereas the Usage capability supports the data requestor. The Access Capability essentially exploits the Negotiation Service to request, view, and possibly update PI that is held by third parties.

The security services (for example, the Open Group taxonomy) are available to all the privacy services and capabilities. The Legal, Regulatory, and Policy Context provide the necessary configuration and parameterization layer.

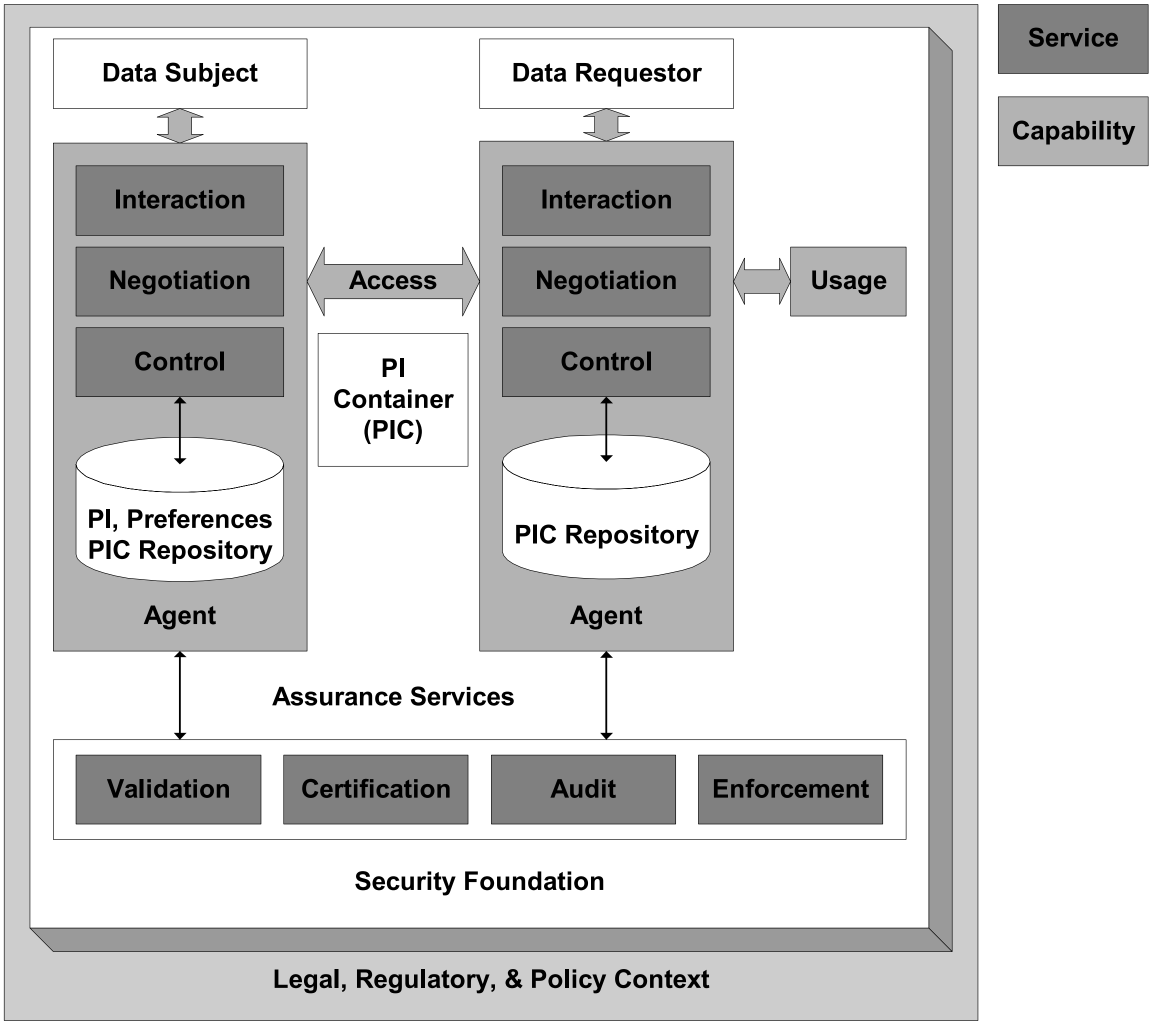

Privacy Framework Services and Capabilities

The original fair information practices can be overlaid on the ISTPA Privacy Framework, showing the operational implementation of the practices. Note that Access is a use case of the Framework, exploiting the Negotiation Service.

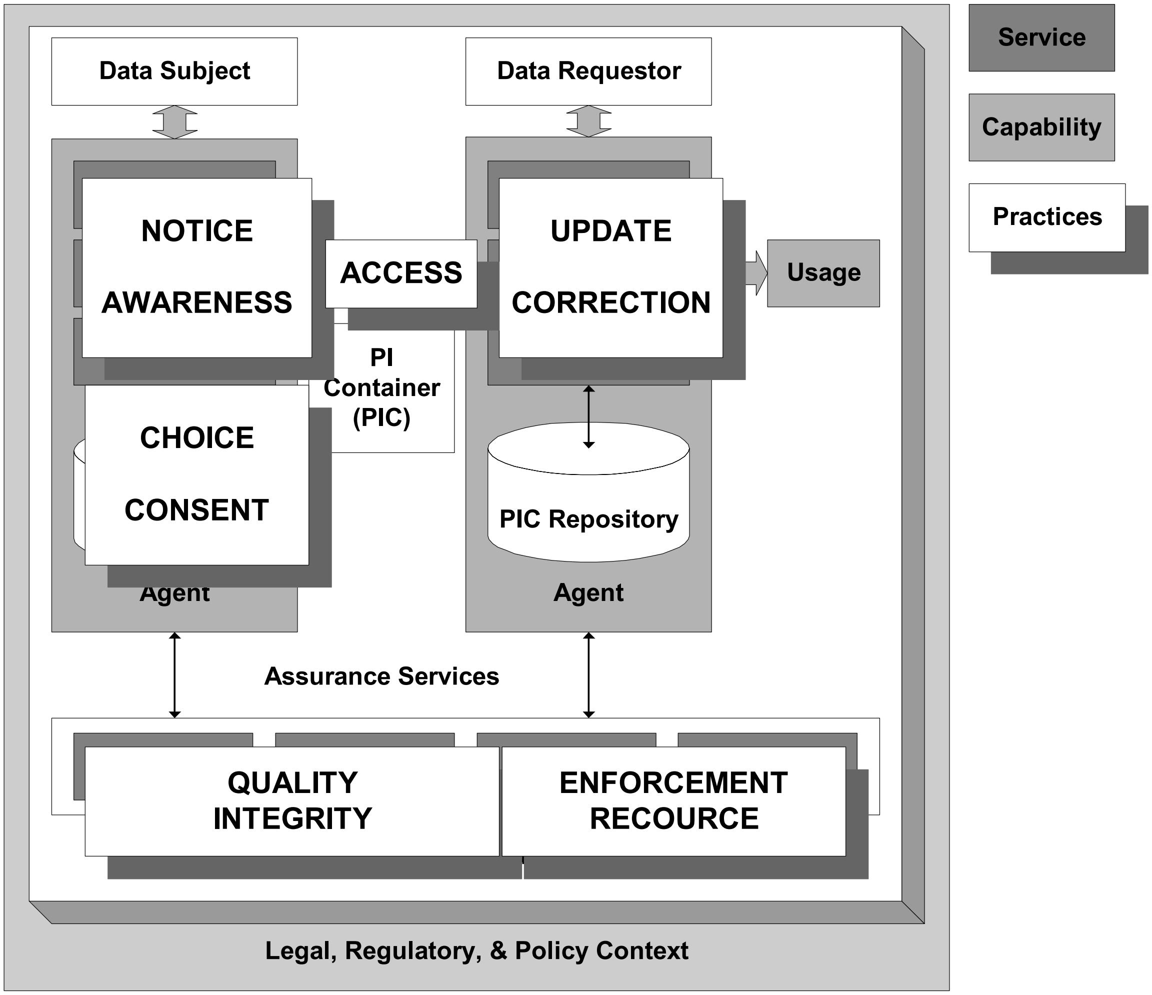

Privacy Framework and Privacy Practices

Description of specific terms

A more complete glossary is located at the end of this document. Listed here, however, is a subset of the key terms that appear frequently in describing the Framework services and capabilities.

Actors

Actors are both individuals and entities such as organizations or computer programs that interact with and invoke the Framework services and capabilities.

Agreement

An agreement is a set of permissions that are negotiated between the data subject and a data requestor, bound together with the associated PI. Context information, such as the intended purpose and any legal or regulatory requirements, may also be part of the agreement.

Data requestor or collection entity

The data requestor is the entity soliciting data from a data subject directly or indirectly from a third party.

Data subject

The data subject is the entity to which the data pertains. Typically this is an individual, but the data subject may also be a corporate entity.

Objects

Objects are the sets of data, including both personal information and programming code, which are acted upon by the services and capabilities.

PI

Personal Information (PI) is any data related to an individual or entity, regardless of whether the subject of the PI is identified.

Permissions

Any activities relating to how the PI is handled that are consistent with the data subject's preferences. Permissions may have been negotiated with a data requestor.

Preferences

Explicit policies and guidelines regarding the handling and use of PI which are agreeable to the data subject.

Service / Capability Interactions

The ISTPA Privacy Framework encompasses all the functions needed to implement the fair information practices, but it is partitioned into subsets of functions that have a logical affinity. The services and capabilities listed above can, in total, support the privacy principles and resolve the fair information practices.

The functions of one service may invoke the functions of another service. In other words, one service may "call" another service (for example, pass data to the other service for subsequent action). In this way, the services (and capabilities) interact in some interconnected sequence to accomplish an overall privacy management task. Use cases will illustrate such interactions and their sequencing as the Framework is used to solve a particular privacy problem. By examining and by solving multiple use cases, the Framework can be tested for applicability and robustness.

Use case scenarios

Use cases involve the sequential interaction of different services and capabilities in order to achieve some desired functionality. A simple scenario is presented here for illustration.

This is the user view of a typical use case. The consumer is browsing, finds a desirable product at a website, is offered a discount in exchange for some PI, but insists on no third-party transfer of the PI. Later, the web-based merchant does re-sell the PI, which raises an alarm that alerts the appropriate authorities.

The consumer wants the confidence that the distributed system will honor the preferences and permissions related to PI that he or she specified when the PI was originally provided.

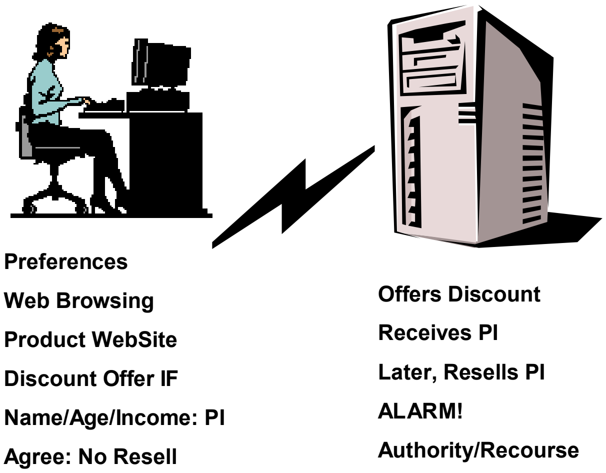

Scenario (Use Case)

The following use case illustrates how the services and capabilities of the ISTPA Privacy Framework interact in this situation.

From the consumer's perspective:

Service / Capability	Data subject's action
Interaction	Consumer determines/configures shopping preferences.
Validation	Checks preferences.
Control	Stores preferences in Repository.
	Consumer browses the web and finds a vendor or retail site with a desirable product.
Negotiation	Consumer views vendor site and considers product discount in which PI (for example: name, age, and income) is required to receive discount offer.
Negotiation	Consumer through Control matches offer with user preferences.
(Agreement reached)	Consumer agrees with the additional "permission" of "No third party transfer of PI"
Control	Consumer's agent stores PI contract, binds permissions & PI, transfers PI container to vendor.
	Product purchase completed.

From the vendor and external auditor perspectives:

Service / Capability	Vendor's action
Interface	Contact established with the consumer.
Negotiation	Vendor offers the customer a discount in exchange for PI. Both sides reach an agreement.
Control	The PI container is transferred and stored under Control with the associated permissions.
	In subsequent processing, the customer's PI is "sold" to a third party, an action not permitted by the agreement..
Usage	System detects violation of permissions (contract)
Usage	An alarm is sent to the Audit Service
Audit	An exception condition results, and the Enforcement Service is contacted.
Enforcement	A notice is sent to the designated external authority.
	The authority initiates recourse actions.

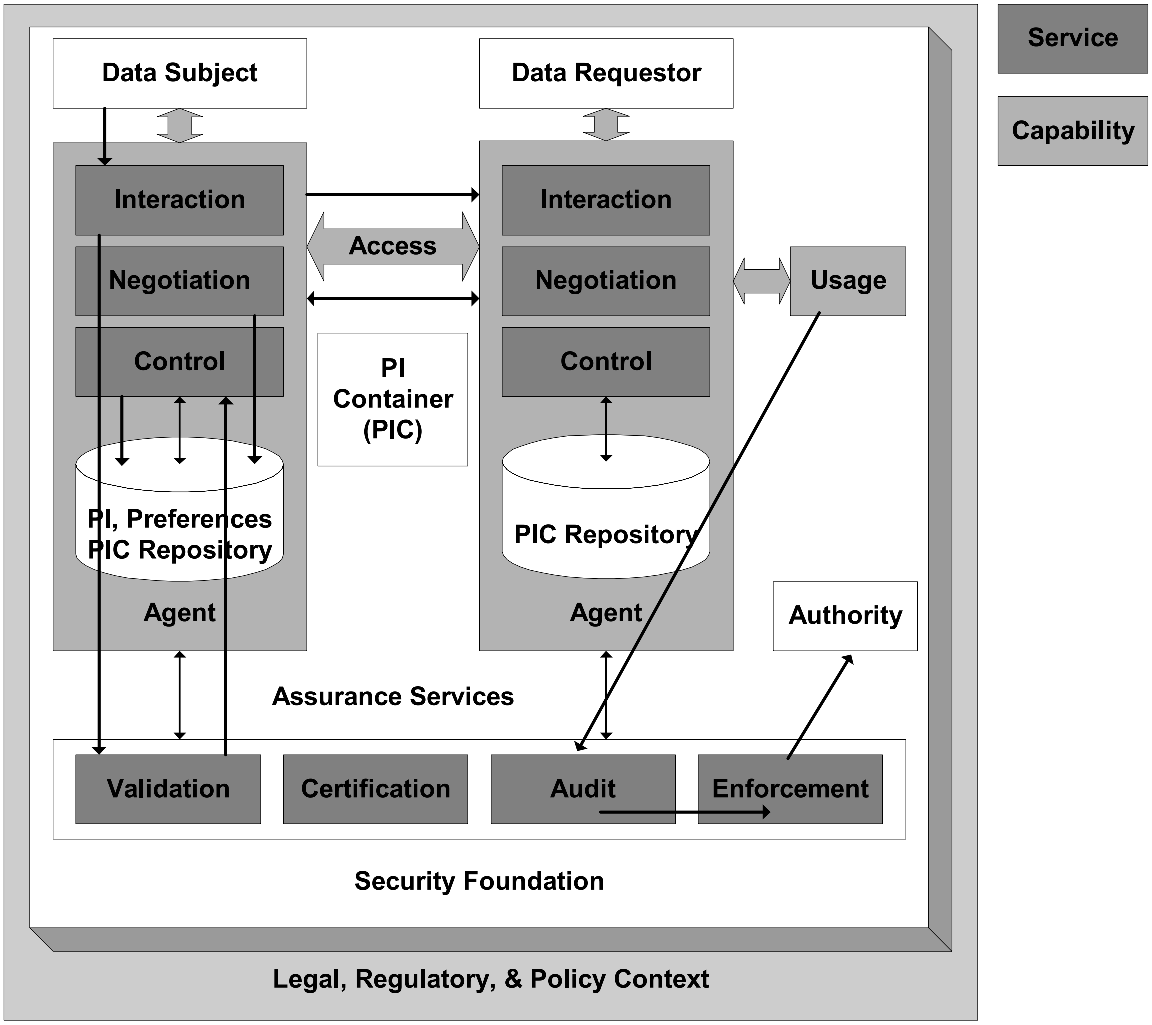

Privacy Framework and Use Case Scenario

The security services (for example, as developed in the Open Group architecture) have only been referenced, but security and security processes are essential to the proper operation of the ISTPA Privacy Framework. Each service and capability in the ISTPA Privacy Framework must be able to call on the appropriate security functions when needed.

Summary

The ISTPA Framework Working Group is charged with developing an administrative, technical, and legal privacy framework within which to provide a functionality that satisfies the privacy principles and fair information practices. A layered analytic approach was adopted, by which the principles/practices were morphed into privacy services and capabilities, which in turn are realized by underlying privacy mechanisms. These services/capabilities call on and integrate with the security services defined by any appropriate security taxonomy (example, Open Group), to satisfy the security principles. The ISTPA Privacy Framework is tested for robustness by considering a variety of use case scenarios.

In the following sections, the ten services and capabilities in the ISTPA Privacy Framework are described at an engineering requirements level. Future technical work will convert these requirements into formal architecture.

The selection of services and capabilities and their resulting behavior in a given context is governed by the particular configuration of privacy policy, operational parameters, and legal and jurisdictional input to the Framework.

The Framework is a “work in progress” that is continually evolving toward more detailed functional descriptions of the services, capabilities, and their interactions. Additional use cases, some with a specific industry orientation or application (e.g., financial services, CRM, eGovernment, location-based services, wireless medicine, mobile commerce and Enhanced 911) will be developed to further test the robustness and applicability of the Framework.

1 Audit Service

1.1 Overview

The Audit Service handles the recording and maintenance of events from other services. It captures the data necessary to ensure compliance with the terms and policies of an agreement and any applicable regulations. The motivating fair information practices are use limitation and accountability.

1.1.1 Privacy regulatory audit requirements

Audit services need to accommodate a host of current and emerging regulatory requirements. These requirements often differ in the breadth and depth of necessary compliance elements that are subject to audit. Whether self-regulatory in nature or statutorily required, several different parties exist: independent auditors, regulatory agents, and internal employees who are conducting self-assessments in order to fulfill audit-reporting requirements. This last group requires trusted, comprehensive, and, in some cases, unannounced means to examine audit material to satisfy audit requirements and perform enforcement or remedy investigations. Similar to security audits, it can be necessary to employ agency or management-approved ethical hacking or surveillance of personnel to attest to the sufficiency, verification, and proper employment of controls. Audit services and mechanisms which attest to the proper and necessary processing of personal information (PI) must support varied audit and investigative techniques to assure any relying parties that PI processing is done according to relevant law and current and evolving market expectations.

Fundamental to most existing privacy regulations, statements of fair information practices, and industry association privacy practice statements, auditable controls must be in place to ensure compliance and effective satisfaction of citizen concerns. It is the security, access, notice, enforcement and remedy principles, often debated in their scope and design, which will drive the scope and detail of audit services and mechanisms.

1.1.2 Audit and certification: other resources

In formulating audit services and mechanisms for the ISTPA Privacy Framework, it is recommended that other works be reviewed and incorporated. One work is the American Institute of CPAs (AICPA) WebTrust Principles and Criteria for Certification Authorities. This work draws on several others (ISO, ANSI, and ABA) in forming an audit and certification foundation for certification authorities (CA). Since many ISTPA members are familiar with CAs and PKI, the AICPA's document is a valuable source of controls, audit, and attestation and certification elements necessary to CA services that can be adapted and incorporated into the ISTPA Privacy Framework.

1.2 Functional description

The key functions of the Audit Service include the following:

Trusted Audit

- Securely trace the processing of all PI including: authoring, access, addition, modification, and erasure.
- Securely trace all privacy preference processing including: authoring, access, addition, modification, and erasure.
- Securely trace all privacy agreement processing including: authoring, access, addition, modification, and erasure.
- Provide strong authentication access by the PI owner or data subject, the regulatory authority, the judicial authority, the PI controller, or the PI requestor.
- Provide a rule-based authorization service that governs all processing by the PI owner or data subject, the regulatory authority, the judicial authority, the PI controller or the PI processor.
- Persistently associate or couple an audit log with the PI. The audit log resides or travels with the PI or sits in a trusted middleware layer surveying all PI object processing, wherever and whenever that processing occurs.
- Transform the association or coupling of audit log(s) across several instantiations or versions of PI objects, and provide consolidated views of PI processing across processing entities (for example, PI owners or data subjects, PI controllers, and PI processors).

Audit Policy Manager

- Match and manage the necessary and applicable privacy principles, practices, and/or regulatory requirements that pertain to PI processing.
- Provide a policy adapter (to interface PI objects and processor entities) that correctly matches the governing privacy principles, practices or regulations that are necessary for compliant processing.
- Provide rule-based intelligence to ascertain which governing rules and necessary audit tracing inventories need to be enabled.

1.3 Actors and objects

The key actors and objects involved in the audit are:

- **Agreement object**, which consists of the completed agreement template after a successful negotiation.
- **Audit object**, which consists of PI processing events.
- **Audit policy manager**, which matches appropriate privacy principles, practices and regulatory requirements that are necessary for PI-protected and compliant processing.
- **Auditor** or **investigative agent**, which require access to audit services.
- **Certification authority**, which issues, suspends, and revokes certificates.
- **Control Service data repository**, where the resulting agreement object is stored.

- **Meta dictionary**, which defines and structures PI elements, audit events and nomenclature.
- **PI controller** or **processor**, which require access to audit services.
- **PI object**, consisting of PI that was either authored by the data subject or the PI owner.
- **PI owner** or **data subject**, which require access to audit services.
- **Preference object**, which details the data subject's privacy preferences governing the PI.
- **Privacy policy,** which details the data controller or processor privacy policy.
- **Regulatory** or **judicial authority**, which require access to audit services.
- **Third parties**, who may wish to process PI or to gain unauthorized access or processing privileges by assuming the identity of legal PI agents or processors.

1.4 Use case scenario

1.4.1 Name

Audit authoring of privacy preference.

1.4.2 List of scenario actors and objects

Audit object, audit policy manager, certification authority, data subject, meta dictionary, PI object, privacy preference object.

1.4.3 Scenario purpose and overview

Purpose: Provide a secure and trusted audit trail of the creation of a data subject's privacy preferences.

Overview: A data subject authors for the first time a set of privacy preferences pertaining to the subject's personal information. Audit services are engaged to create a secure and authentic recording of this event, and to securely store that event for processing by the regulatory authorities or other authorized agents.

1.4.4 Actor or action and system response

Actor / Action	System Response
1. The data subject opens a privacy preference-authoring application.	
2. The data subject enters necessary access codes to gain access to PI and associated privacy preference(s).	3. The system challenges and authenticates the data subject.

Continued…

Actor / Action	System Response
	4. The system opens an audit object and records the PI and preference access event.
5. The data subject selects one or more PI elements to bind with a privacy preference.	
	6. The audit policy manager determines the necessary PI privacy preference rule to drive the Audit Service event recording level.
7. The data subject binds the PI with selected privacy preference(s).	8. The Audit Service records the PI to a privacy-preference binding event according to the governing privacy preference rule.
9. The data subject closes the privacy preference authoring application.	10. The Audit Service records closure event.
	11. The Audit Service secures an audit object and sets the necessary access rules.
	12. The Audit Service indicates to the System that the application can be closed.

1.5 For further consideration

Audit capabilities, events, and nomenclature need to be defined and managed. Many audit capabilities will record a host of events necessary to ascertain the ongoing trust, security and privacy of PI processing in a network computing environment. Leveraging an audit framework already in use by industry will be critical.

In an environment of evolving and diverse data protection regulations, privacy policies, and security threats and vulnerabilities, it will be critical to develop an adaptive audit control mechanism. Efforts to harmonize, reconcile or standardize legal and regulatory requirements for data protection are underway. It will be important to keep track of these efforts by creating a “rule-repository” that collects and administers a change-control capability that continuously assesses and adapts the controller’s systems to governing data protection requirements.

Certification services will play a significant role in audit services. The issuance, revocation and retirement of credentials necessary to control access to audit services and resultant objects will be essential.

2 Certification Service

2.1 Overview

The Certification Service supports the management and affirmation of credentials of any responsible party involved in processing personal information (PI). Security and trust necessitate services that certify or attest to actor compliance and trustworthiness. The Certification Service relies on audit services that account for PI processing. The credentials of those actors are issued and maintained in accordance with applicable data protection regulation and management standards.

Whether self-regulatory in nature or statutorily required, several certification (e.g., seal) programs exist, each with differing criteria as to what constitutes trust, security and privacy compliance. A trusted certification manager will be necessary to compile, ascertain and mediate the differences in jurisdictions and ongoing changes in what is or is not a necessary PI processing requirement. Also, market forces will dictate particular policies, practices and assurance criteria, with accreditation of actors incorporating those market-driven criteria.

Differences of opinion as to what these criteria should be become apparent as businesses analyze existing privacy regulatory requirements (COPPA, HIPAA, GLB, EU Data Protection Directive), statements of fair information practices (OECD, CSA, FTC, Article 29 Committee), and industry association privacy practice statements (DMA, OPA, Global Dialogue). Consistent across all of these, however, is the idea of a notification mechanism to demonstrate compliance. A basic service that notifies relying parties of an actor's level, type, or status of compliance, is necessary.

2.2 Functional description

The Certification Service is comprised of mechanisms that affirm the credentials of any party involved in controlling or processing a PI object or PI transaction. The key functions of the Certification Service include the following:

- Ascertain relevant privacy principles and/or regulatory requirements, and assemble necessary criteria for certification.
- Maintain a trusted certificate display and notification service in order to alert relying parties, subscribers, and authorities of the data controller's or the processor's certification status.
- Compile certificate issuance and certificate revocation histories for regulatory and public examination.
- Maintain audit and attestation workflow progress for PI controllers and processors under certification examination.
- Assign independent auditors or agents for any monitoring and investigative procedures that are necessary for certificate maintenance. Provide a certification or compliance warning system that alerts relying parties of lapses, violations, or potential for adverse affects on the data controller or processor's certification status.

2.3 Actors and objects

The key actors and objects involved in the Certification Service are:

- **Agreement object**, which consists of the completed agreement template after a successful negotiation.
- **Audit object**, which consists of PI processing events.
- **Audit policy manager**, which matches appropriate privacy principles, practices, and regulatory requirements necessary for PI protection and for compliant processing.
- **Auditor** or **investigative agent**, which requires access to the audit services.
- **Certification authority**, which issues, suspends, and revokes certificates.
- **Control Service data repository**, where the resulting agreement object is stored.
- **Meta dictionary**, which defines and structures PI elements, audit events and nomenclature.
- **Personal Information (PI) controller** or **processor**, which requires access to the audit services.
- **PI object**, which consists of the PI that is authored either by the data subject or by the entity that owns the PI.
- **PI owner** or **data subject**, who requires access to the audit services.
- **Preference object**, which details the data subject's privacy preferences governing the PI.
- **Process certificate object**, which binds the necessary credentials to the data controller or processor and points to issuing authorities' examination report.
- **Process certificate manager**, which ascertains relevant privacy principles and/or regulatory requirements, assembles the necessary criteria for process certification, and manages the certificate life cycle.
- **Privacy policy**, which details the data controller or processor privacy policy.
- **Regulatory authority** or **judicial authority**, which requires access to the audit services.
- **Third parties**, who may either wish to process the PI legally or possibly gain unauthorized access or processing privileges by assuming identity of legal PI agents or processors.

2.4 Use case scenario

2.4.1 Name

Relying on a process certificate.

2.4.2 List of scenario actors and objects

Audit policy manager, certification authority, data subject, PI object, process certification manager, process certificate object

2.4.3 Scenario purpose and overview

Purpose: Provide a secure and trusted certification process for relying parties to ascertain whether to engage a particular data collector or controller.

Overview: A data subject checks the credentials of a data controller before exchanging PI with the data controller. The certification services are engaged to create a secure and authentic display of a credential and then to record the reliance by a data subject.

2.4.4 Actor or action and system response

Actor / Action	System Response
1. The data subject queries a data controller’s privacy policy and examines the certification credential.	
2. The data subject enters a data controller’s website and accesses its privacy statement.	
3. The certification manager records an inquiry to a certified data controller.	4. The system opens an audit object and records the relying party’s inquiry.
5. The data subject activates a certificate status check.	
	6. The certification manager opens a certificate object and securely displays a credential status.
	7. The certification manager accesses a revocation list server and performs a revocation check.
	8. The certification manager alerts the regulatory authority of the relying party’s inquiry and records a certificate object status report to the data subject.
9. Contingent upon a certificate status, the data subject continues interaction with data controller, or ceases interaction.	
	10. The certification manager notifies the Audit Service and closes the audit object.

2.5 For further consideration

Harmonization of disparate data protection laws and necessary criteria for certification authorities will drive much of what is being verified.

Audit and attestation service providers already have efforts underway and should be consulted.

PKI is still evolving, with regulatory, market forces, and technology forging a working approach to certificate life cycles and mechanisms for protecting and relying upon these digital credentials.

The ongoing interplay of industry, government, media, and privacy advocates will affect the trustworthiness of seal programs and other credentials, and whether they are acceptable to the users.

3 Control Service

3.1 Overview

The Control Service encompasses a number of functions. These functions work together to ensure that the fair information practices operate according to prescribed privacy policy on personal information (PI), which is maintained and manipulated by either a data collection controller or a data processing entity.

The Control Service touches on all aspects of the privacy principles, but is centrally concerned with the principles of use limitation, accountability, security, and verification.

Within the scope of the Control Service are the functions that ensure compliance with the fair information practices of choice and consent, by ensuring the data is processed in accordance with the established agreements. The Control Service also ensures compliance with the fair information practices of enforcement and recourse, by providing clear audit trails that can expose misuse. Notice and awareness are supported by allowing data subjects to determine how their PI either will be or has been used.

3.2 Functional description

The Control Service is comprised of the functions that enforce the agreements, policies, and regulatory requirements that are applicable to a given element of PI. Key functions of the Control Service include the following (with explanation):

- **Surround PI in repositories that are under the control of the data processing entity.**
- **Selectively allow PI to flow in and out of the repositories, but such flows are subject to conditional approval by the Control Service.** Data may enter a repository from a number of sources, collectively referred to as data providers. Typically, data providers may represent some automated information-gathering process that interacts with the Data Subject via the Internet. However, not all of the data provided to a repository will necessarily come from automated collection channels. Data entering the repository may be sourced in traditional data channels, such as paper forms or telephone conversations.
- **Link the PI to a data usage agreement.** Regardless of the channel that provides the information to the repository, the Framework specifies that PI that is placed into a repository by a data provider be linked or bound to a data usage agreement. In cases where PI enters a repository without an associated data usage agreement, a separate task of linking the PI to a data usage agreement will be invoked.
- **Interact directly with internal requestors.** Internal requestors are defined as those requestors that are implicitly or explicitly bound to the terms of data usage agreements. Typically, this means employees and agents of the data processing entity. In some cases, this may also imply third parties that are contractually bound to abide by the terms of the data processing entity. From the perspective of the data usage agreements, internal requestors must act as the data processing entity. Related requestors that are bound by the terms of the data usage agreements, but which are not the data processing entity, will generally be considered third parties for the purpose of determining access permissions and should not be considered internal requestors. The specific language of

the data usage agreement must clarify exactly what latitude a data processing entity may have in defining close affiliates and subsidiaries as internal requestors.

- **Interact indirectly with external requestors.** External requestors are entities that are seeking PI from the repository. They are not implicitly bound by the data usage agreements, nor are they considered part of the data processing entity. The Control Service will grant or prevent external requestors access to PI under the terms of the data usage agreement. Such access is predicated on the development of a new data usage agreement between the data processing entity and the external requestor. The terms of this data usage agreement cannot be inconsistent with the terms of the data usage agreement that exists between the data processing entity and the data subject.

  Typically, external requestors are obligated to abide by the terms of the initial data usage agreement and are also obligated to abide by additional constraints imposed by the data processing entity. The external requestor will also agree to maintain the binding between the PI and the governing data usage agreements, as a condition for receiving the PI.

- **Receive specifically defined rules from the data processing entity policy manager.** Operation of the Control Service is governed by laws, regulations, data processing entity policies, and data usage agreements. The data processing entity translates law, regulation and data processing entity policies into specifically defined rules that are then provided to the Control Service. Policy managers for the data processing entities are responsible for maintaining, revising, and transmitting the rules to the Control Service.

- **Provide data for audit/logging regarding any PI transactions that are under the purview of the Control Service.** Additionally, the Control Service provides audit data that relates to the creation or changes in rules provided by the data processing entity policy manager. The audit policy settings determine which specific data has to be logged.

3.2.1 The Control Service receives the following inputs:

- Audit settings, which define the audit data to be generated by the Control Service during its operation.
- Control service rules, which are provided by the policy manager and, in concert with the data usage agreements, govern the behavior of the Control Service.
- Data usage agreements.
- Personal Information (PI).
- Requests for access to stored PI. Requests for PI are accompanied by enough detail to cover the nature of the requestor and the intended usage of the PI. This detail then allows the Control Service to determine if the request is allowable under the control service rules and the data usage agreements.
- Requests for the storage of PI. These requests are accompanied by a data usage agreement, or a reference to a data usage agreement that applies to the PI. PI without a data usage agreement may be accepted on a temporary basis, pending the generation of a governing data usage agreement. Additionally, the PI provided will be appropriately categorized to allow the Control Service to determine if the storage of such data is consistent with law, regulation, and data processing entity policy as encoded into the Control Service rules.

3.2.2 The Control Service manages the following outputs:

- Audit Data.
- Personal Information (PI).
- Rejection of Requests: The Control Service may reject requests for access to PI and/or requests to store the PI.
- Processing.

The Control Service is essentially a rules processing engine. It accepts a variety of requests relating to the repository that is controlled by the service, and determines if the request is to be honored.

The decision-making process within the Control Service compares the general rules that relate to the request (storage, retrieval, binding, etc) with the specifics of the request. The Control Service also examines the data usage agreement associated with each PI element that is covered by the request. If no general rule prohibits the servicing of the request, and if the data usage agreement covering the affected PI element allows the request, then the request will be honored.

The sophistication of the rules engine is not addressed within the Framework. Implementers are free to innovate in the development of advanced engines for this service, but basic operation requires that the Control Service enforce the legal obligations of the data processing entity regarding the access and usage of PI.

3.3 Actors and objects

The key actors and objects involved in the Control Service are:

- **External requestors.**
- **Internal requestors.**
- **Personal information (PI).**
- **Privacy policies.**

A central concept of the Privacy Framework is the persistent linkage and binding between data usage agreements and PI that permit a data processing entity to use PI in specific ways. The PI element tied to an agreement becomes a PI object; the behavior (i.e., agreements) and data are now bound together. This section will continue to use PI to refer to both personal information and personal information objects. This object is also referred to as the PI container.

The Control Service interacts with numerous entities internal to the data processing entity. Acting at the repository boundary, the Control Service activates each time the repository is accessed.

3.4 Use case scenario

3.4.1 Name

Advertising solicitation

3.4.2 List of scenario actors and objects

Data requestor, personal information

3.4.3 Scenario purpose and overview

Purpose: Respect customer preferences and permissions.

Overview: A company creates a database with customer names and addresses that it has collected from its website. In the database, each customer record is linked to a record that contains the allowable uses that the customer agreed to at the time the information was collected. Included in this record is the customer's response (via a check box) to the question: "May we use your address to send you information about future products?"

The company now has a new product that it wants to advertise to its customers. A marketing manager in the company sends a request to the database for mailing labels. This request identifies what data is being requested, who is the requestor, and what is the intended use of the data.

At this point, the Control Service receives the request and the associated request information. The Control Service evaluates the request and determines that the request does not violate any of the general rules relating to the database. The query is allowed to proceed, retrieving both the customer records and the associated customer permission records (called data usage agreements). For each customer, the setting in the data usage agreement (from the above checkbox) is examined, and only the records for the customers who provided this usage permission are included in the output. Concurrently, the Control Service generates audit data relating to this request. Individual customer records might also be updated to reflect the fact that the information was provided in response to this request.

At the end of the process, the marketing manager receives a list of names and addresses for customers that have agreed in advance to allow such usage. The manager is obligated to use the data only for the purpose indicated in the request. To further ensure this constraint, we assume that the list could contain a permitted use statement advising the manager that only the stated usage of the data is permitted.

3.4.4 Actor or action and system response

Actor / Action	System Response
	1. The system has already collected the data subjects' PI plus permissions in a secure repository.
2. The data requestor asks for mailing list, accompanied by an intended use statement.	3. The system evaluates the request; no violations.
	4. The system extracts only the records for which the permissions match the data usage agreement.
	5. The system sends the list to the data requestor.

Continued…

Actor / Action	System Response
	6. The system creates an audit entry.
7. The data requestor receives the list with a reminder about its intended use.	

3.5 For further consideration

Managing migration into legacy systems can be a challenge. The Control Service is central and fundamental to the robust protection and proper handling of PI. The challenge is migrating the Control Service into an existing information system.

Consistently enforcing permissions that are under the control of third parties can be a challenge. After permissions are negotiated and bound to a data subject's PI, the PI container is forwarded to the entity that is requesting the PI. Potentially, the PI and container are also shared with third parties. There is a potential for the third parties to misinterpret or to not consistently enforce the permissions associated with the PI. As the Framework model is adopted, due diligence, branding, and corporate reputation will require that requestors and third parties faithfully implement the Control Service, which will enforce the permissions.

4 Enforcement Service

4.1 Overview

The Enforcement Service initiates response actions when a data collection entity does not conform to the terms or policies of an agreement or the applicable regulations. The motivation for the Enforcement Service is the fair information practices principle of accountability. Enforcement also includes the Recourse function, discussed below, which provides recourse for data subjects when their PI is being used differently from the original agreement. Examples are given in the context of financial institutions, but similar monitoring, audit, and redress options exist in other industry sectors.

Public Sector

Ongoing monitoring is prescribed by regulation and published standards. Examples include U.S. financial services oversight by the Office of the Comptroller of the Currency (OCC), and federal and state banking regulators. In this model, regulators document out-of-compliance situations and as a consequence can either cause corrective action and/or punish violations. The primary mechanisms typically include fines or reputation consequences.

Exception based monitoring also may be prescribed by legislation, but may require a triggering event to activate. Examples include the Federal Trade Commission's monitoring of third party firms (those that have no formal financial services charter) who provide financial services as defined in Title 5 of the Gramm-Leach-Bliley (GLB) legislation. Other examples include standard FTC deceptive-trade practices investigations for retailers and other service providers. In an exception environment, regulators act upon out-of-compliance evidence that may come from a number of sources, such as customers, third parties, or routine regulatory sweeps. Depending upon the nature of the violations, consequences can include a wide variety of actions.

Private Sector

Enforcement in the private sector typically includes ongoing processes associated with basic, routine business practice. Primary examples are regular third-party audits and internal compliance programs. Here, feedback loops provide direction to internal actors, up to the board of directors, who ensure corrective action. Public consequences can include audit exceptions with resulting reputation and valuation consequences.

Voluntary enforcement programs enable service providers to subscribe to a set of guidelines to which they agree to adhere, and – most relevantly – these programs typically contain procedures for self-certification, third-party certification, or a combination of both. Generally, self-certification is considered an adequate base in the absence of indicators suggesting abuse, in which case a third-party review is mandated. There are a wide variety of such programs. Feedback loops in this environment lead to consequences imposed by the third-party program sponsor, and might include withdrawing the organization's seal of approval. As in the example above, these programs rely on the market itself to punish out-of-compliance behaviors.

Laissez-faire market mechanisms may also provide meaningful enforcement mechanisms, especially in online environments where information can flow quickly in real time. In this model, aberrant behavior as defined by the market itself is publicized through normal communication channels and can lead to virtually instant reputation and valuation consequences, causing permanent damage to the business.

4.1.1 Recourse Function

Public Sector

Referral mechanisms, which range from simple publicly accessible databases to automated services, help determine which actor has jurisdiction in a given suspected out-of-compliance situation. Note that these mechanisms can work at several levels: (1) to identify the actor within a business who might most appropriately respond to a customer complaint; (2) identify the regulatory authority with oversight and redress authority; (3) identify the specific individual or office within the regulatory authority who is the best point of contact to resolve the issue.

Registration mechanisms enable customers to formally lodge a complaint, either within a business or with a supervisory authority.

Evaluation mechanisms determine the validity of a complaint. Some mechanisms are based on ongoing process evaluation associated with regulatory oversight; others on exception investigations prompted by an individual incident but are still performed by regulatory entities with appropriate jurisdiction.

Redress vehicles determine the nature of “make good” efforts and ensure they are carried out. These efforts might take several forms, from simply fixing the outstanding issue, to enforcing a change in ongoing business process, to determining and applying a wide range of penalties.

Private Sector

The same mechanisms noted above apply in a self-regulatory environment, except that a third-party agent would play the role of a regulator. Typically, the range of penalties (beyond reputation and valuation consequences) in such an environment might be more limited than in the public sector examples noted above.

One example is an online dispute resolution vehicle, such as NovaForum.com. NovaForum offers alternative dispute resolution techniques, including mediation, arbitration, and neutral evaluation (also known as non-binding arbitration). Companies sign up for this service on a subscription basis. Dispute resolution is targeted for 72 hours from the time a customer files a complaint. The service can be used in both B2B and B2C marketplaces.

4.2 Functional description

The Enforcement Service comprises both the Enforcement Service and the Recourse function. The primary processes that make up the Enforcement Service are feedback loops, comparing expected performance with observed performance, and mechanisms for imposing consequences in situations where the observed performance is out of compliance. The primary processes that comprise the Recourse Function are customer identification of the relevant authority, initiation of complaint by either the customer, a third party, or the regulator; testing the complaint’s validity, and where complaints are valid, determination of both specific redress and a fulfillment process.

The key functions of the Enforcement Service include the following:

- Provide data subjects a means to report alleged violations of privacy policy to enforcement and advocacy agencies and to certification and judicial authorities.
- Provide referral mechanisms to aid data subjects in determining which enforcement or remedy actor has jurisdiction in a given suspected violation.
- Provide data subjects, data controllers, processors, regulatory authorities and enforcement agencies access to complaint/remedy databases.

- Provide regulatory authorities and enforcement agencies access to audit services.
- Provide data subjects, data controllers, and processors access to dispute resolution services.
- Provide regulatory authorities and enforcement agencies the means to compile investigative findings, as well as access to the forensic audit logs and investigative findings of auditors, investigative agents, and certification authorities.
- Provide regulatory authorities and enforcement agencies the means to impose remedies, or corrective actions, and to secure relief on behalf of data subjects and data controllers.
- Provide regulatory authorities and enforcement agencies the means to request updates for the credentials and certification status of data controllers and processors.
- Provide the means for data subjects, controllers, processors, certification authorities, regulatory authorities, enforcement agencies, and judicially authorities to exchange investigative findings and compile authoritative records.
- Provide regulatory authorities and enforcement agencies the means to track and prosecute false claims and disruptive acts intended to defraud or damage data controllers and processors.

4.3 Actors and objects

The primary actors and objects involved in the Enforcement Service are dispute resolution authorities. These entities evaluate situations where an out-of-compliance situation may exist, and, where necessary, impose corrective action including redress.

The key actors and objects involved in the Enforcement Service are:

- **Agreement object**, which consists of the completed agreement template after a successful negotiation.
- **Audit object**, which consists of PI processing events.
- **Audit policy manager,** which matches the appropriate privacy principles, practices, and regulatory requirements that are necessary for PI protection and for compliant processing.
- **Certification authority,** which issues, suspends, and revokes certificates.
- **Control service data repository**, where the resulting agreement object is stored.
- **Personal Information (PI) controller** or **processor**, which requires access to the audit services.
- **PI object**, which consists of the PI that is authored either by the data subject or by the entity that owns the PI.
- **PI owner** or **data subject**, who require access to the audit services.
- **Preference object**, which details the data subject's privacy preferences governing the PI.
- **Privacy policy,** detailing the data controller or processor privacy policy.
- **Process certificate object**, which binds the necessary credentials to the data controller or processor, and points to the issuing authorities' examination report.

- **Process certificate manager**, which ascertains the relevant privacy principles and/or regulatory requirements, assembles the necessary criteria for process certification, and manages the certificate life cycle.
- **Regulatory authority** or **judicial authority**, which requires access to the audit services.
- **Third parties,** who may either wish to process the PI legally or possibly gain unauthorized access or processing privileges by assuming identity of legal PI agents or processors.

4.4 Use case scenario

4.4.1 Name

Privacy violation review by an enforcement agent

4.4.2 List of scenario actors and objects

Agreement, audit object, audit policy manager, certification authority, data subject, enforcement agent, meta dictionary, PI object, privacy policy, privacy preference object.

4.4.3 Scenario purpose and overview

Purpose: Enforcement agent accesses and reviews consumer complaint.

Overview: Consumer files a privacy violation complaint with a privacy seal program agency. The privacy seal program agency reports the complaint to the governing enforcement agency. The enforcement agency assigns an enforcement agent the tasks of reviewing the complaint and conducting an investigation.

4.4.4 Actor or action and system response

Actor / Action	System Response
1. The enforcement agent accesses the complaint database and requests the complaint and its associated investigative objects.	2. The system identifies and authorizes access to requested investigative objects.
	3. The system displays the complaint, PI object, privacy policy, privacy preference object, and audit object.
4. The enforcement agent reviews the complaint and its associated investigative objects.	

Continued…

Actor / Action	System Response
5. The enforcement agent consults audit policy manager to assess any governing regulatory requirements and associated penalties and remedies.	
6. The enforcement agent requests the data controller's certification histories and any records of other complaints by data subject during the last three years.	7. Certification history and data subject complaint files are provided.
8. The enforcement agent closes the session and sets a date for management review.	9. The audit manager records the enforcement agent's request and delivery of requested investigative objects.
	10. The system ends the session.

4.5 For further consideration

How are audit, enforcement, and recourse processes scaled so that they will be economically feasible to operate on a continuing basis across a wide range of businesses? The outline above is based primarily on financial services industry models, and it defines a set of processes that are most appropriate to a highly regulated, closely monitored industry. The outline is not meant to imply that a similarly broad set of processes, rigorously imposed, is appropriate for other industries. Indeed, a major concern is that, in a laissez-faire environment where the costs of extra compliance and recourse burdens are not mandatory, too much additional infrastructure could prove unsustainable.

Are voluntary arbitration services workable on a large scale?

How do voluntary arbitration services compare to the traditional Better Business Bureau approach?

To what extent are Audit, Certification, and Enforcement Services integrated? Under what circumstances is such integration constructive, or not?

5 Interaction Service

5.1 Overview

The fair information principles and practices that are reflected in the Interaction Service include notice and awareness, choice and consent, and collection limitations. Interaction also includes openness, to the extent that a generalized interface is required for raw information presentation and handoff. The concepts embodied in the Interaction Service are not expressly related to privacy, but the service is defined to emphasize the wide range of interactions between the Framework and all entities outside the Framework.

The Interaction Service handles the external interfacing, whether to a human, a computer, or some external automatic input/output process. The internal interfaces between Framework components and other services and capabilities are currently defined informally, but could be defined as formal interfaces as part of a technical Framework architecture in the future.

5.2 Functional description

The key functions of the Interaction Service include the following:

- Handles presentation of data between entities outside the Framework structure and components inside the Framework, such as input of the data subject's personal information (PI), preferences, and actions, as well as confirmation of actions.
- Comprises the external interface to the agent in cases wherein the data subject is represented by an agent. The Interaction Service provides a generalized interface and presentation function.
- Includes input/output elements such as PI, preferences, actions (plus confirmation), permissions, and agreements. Appropriate elements can be either pushed or pulled by the Service; that is, either requested or offered. Mechanisms include methods for data representation, multi-modal I/O (on and off line), communications interfaces, storage of raw data, traditional presentation services (e.g., GUIs), as well as external machine and automation interfaces. Notice and awareness information to be presented to the subject are handled by the Interaction Service.

5.3 Actors and objects

The key actors and objects involved in the Interaction Service are:

- **Data collection entity**, which acts as a requestor of data.
- **Data subject**, which is the source of the PI.
- **Extra-Framework entities**, which function as correspondents with the Framework.
- **Data objects**, which are any data that can be passed into and out of the Framework.

5.4 Use case scenario

5.4.1 Name

Notice and Awareness

5.4.2 List of scenario actors and objects

Data collection entity, data subject entity, relevant PI.

5.4.3 Scenario purpose and overview

Purpose: Provide notice/awareness to the consumer.

Overview: The Interaction Service is used to provide notice and awareness information to the subject consumer, as well as to provide the user interface for passing PI and consumer preferences into the Framework.

5.4.4 Actor or action and system response

Actor / Action	System Response
1. The data subject provides personal information and preferences to a PI configuration tool.	2. The system receives, validates, and securely stores PI and preferences.
3. The data subject browses to a merchant website, referred to now as the data requestor.	
4. The Data requestor presents a request for PI from the data subject, together with the purpose for the data.	5. Notice and explicit awareness of the PI collection request are presented to the data subject by the Interaction Service.
6. The data subject decides whether to reject the request or to enter into further dialog with the merchant website.	

5.5 For further consideration

Provide a consistent, user-friendly, and intuitive experience for the user/consumer. The Framework will be embedded in a variety of environments, each with its particular challenges for external interfacing. For example, wireless devices have memory, computation, and screen size constraints not common to larger computing and communications systems. Creative techniques are needed to preserve the clarity and simplicity of privacy management in such contexts.

6 Negotiation Service

6.1 Overview

The Negotiation Service is motivated by the fair information practices of choice and consent. In essence, these principles hold that an individual about whom PI is collected should understand the purposes for which the data will be used and have an opportunity to provide or deny consent or optionally conduct a negotiation. Through the Negotiation Service, individuals and data controllers can negotiate data collection, usage, and privacy protection terms and conditions. Due to the emerging array of statutory privacy protections, evolving variety of information-for-value offers, and growing consumer awareness and concern about privacy, it is anticipated that template agreements, which employ base-line regulatory compliant and market-acceptable terms, will emerge to simplify and standardize such negotiations.

6.2 Functional description

The key functions and elements of the Negotiation Service include the following:

- The means for the data collection entity and a data subject to negotiate and execute a data protection agreement to which the data subject is a willing and informed participant.
- An agreement that defines the purpose(s) of the data collection so that the user is making an informed decision.
- An agreement that defines what PI is being requested and whether it is required or optional so that the user is making an informed decision.
- An agreement that defines any additional permissions that are being requested, for what purpose(s), for which PI, and whether these permissions are required or optional. While it can be efficient for a vendor to request data for multiple purposes at one time, this needs to be clearly communicated so a data subject can provide informed consent.
- An agreement that directly incorporates or provides hyperlinks to the full text of the policies governing the agreement; i.e., all relevant privacy and security policies under which this data collection is taking place. Ideally, any key points of these policies that are germane to a particular agreement are summarized for the individual.
- An agreement that provides information regarding the effective duration of the agreement so that the user is making an informed decision.
- An agreement that provides information regarding access to the agreement for future review or editing. This provides a bridge to the Access Capability.
- An agreement that provides information regarding whether and to what extent the agreement can be modified. This is necessary for informed consent, and for governing changes that can be made under the Access Capability.
- A document request facility that provides the data subject with the ability to obtain a copy of the data exchange agreement. This is a practice that enhances trust and accountability and facilitates access at a later date.

6.2.1 Processing and mechanisms

The Negotiation Service can be thought of as a mediator that sits between the Agent Capability representing the data subject and the Agent Capability representing the data collection entity. Thus, when the Negotiation Service is performed over a data network, the processing steps are as follows:

- A data subject event triggers the Agent Capability to request either a proposal or conditional agreement object from the Negotiation Service.
- The Negotiation Service requests the proposal or conditional agreement object from the proposal object repository, which may be provided by the Control Service.
- The proposal object repository returns the proposal or conditional agreement object.
- The Negotiation Service returns the proposal or conditional agreement object to the Agent Capability.
- The Agent Capability obtains the necessary inputs from the data subject.
- The Agent Capability returns to the Negotiation Service either: a) a modified proposal object, b) a conditional agreement object, or c) an agreement object (if the object that was sent to the Agent Capability was a conditional agreement object and the data subject agreed).
- The Negotiation Service processes the returned object according to the negotiation rules established by the data collection entity. This can include approving the duration and permissions, validating either the authentication credentials or the data using the Validation Service, or other conditions established by the data collection entity.
- If the outcome requires additional negotiation, processing repeats as necessary. If the outcome is an agreement object, the Negotiation Service returns an acknowledgement to the Agent Capability and posts the agreement object to the Control Service.
- The Control Service acknowledges the Negotiation Service.
- The Negotiation Service sends a final confirmation to the Agent Capability, completing the negotiation.

Note that each action by the Negotiation Service may be logged for audit purposes.

6.2.2 Proposal and agreement state and transformation

Once one party, subject only to agreement by the other party, has agreed upon the terms of a proposal object, it changes state to a conditional agreement object. Once the other party has agreed to the terms of a conditional agreement, it changes state to an agreement object. The only difference between a proposal object and an agreement object is that data requests and permission requests have been turned into data references. These reflect the specific data instances covered by the agreement and the permissions reflecting the specific permissions that are granted or denied by the data subject.

The PI container is used to bind the PI to any agreed-to permissions. The container also contains the intended use, policies, and conditions related to the permissions that were granted, so that the proper context is captured and preserved. Credentials, including identity and digital signatures, are also contained in the container.

Permissions are classified as modifiable or non-modifiable. After negotiation of the original agreement, the data subject can choose to change a modifiable permission without re-negotiation of the agreement. However, a non-modifiable permission may only be changed via re-negotiation of the agreement.

6.3 Actors and objects

The key actors and objects involved in the Negotiation Service:

- **Agreement duration**, which is the period of time the agreement will remain in effect.
- **Agreement object**, which results from a successful data collection negotiation.
- **Agreement PI object requests**, optional or required, which specify the PI elements that have been requested by a proposal.
- **Agent**, which represents respectively the data subject and the data collection entity entering the negotiation.
- **Authentication credentials**, which are used to validate the identity of the parties.
- **Conditional agreement object**, which represents a proposal object whose terms at least one party has agreed to, subject only to approval of the same terms by the other party. Note that a proposal object that the data collection entity will accept a valid response for is already a conditional agreement object.
- **Control service data repository**, where the resulting agreement object is stored.
- **Data collection purpose(s),** which are the stated reason(s) for the collection of PI from a data subject.
- **Optional permission requests**, optional or required, which are requests to use the PI that is covered by an agreement for a purpose other than the primary purpose(s) defined in the purpose element above. A permission contains the same elements as an agreement, except policies and authentication credentials. These are defined by the containing agreement. It can be considered “an agreement within an agreement.”
- **PI controller** or **processor,** which controls or processes PI.
- **PI object**, which consists of the PI authored by the data subject or PI owner.
- **PI owner** or **data subject**, the individual or entity who owns and authors the PI.
- **Policy references**, which are used to assert the privacy, security, and other relevant policies of the data collection entity covering this agreement.
- **Proposal object**, which contains authentication credentials, policy preferences, data use purposes, agreement duration, requested PI objects and permissions.
- **Proposal object repository**, that stores proposal or conditional agreement objects on behalf of the data collection entity. Note that this may be managed by the Control Service.

6.4 Use case scenario

6.4.1 Name

Consumer and e-commerce site agreement negotiation

6.4.2 List of scenario actors and objects

Agent, data subject, control service data repository, conditional agreement object, optional permission request, PI object, policy references, proposal object.

6.4.3 Scenario purpose and overview

Purpose: Demonstrate negotiation service in an e-commerce transaction.

Overview: A consumer wishes to complete an e-commerce transaction and is presented with a conditional agreement with several optional permissions. The consumer views and selects the permissions they deem appropriate and responds to the data controller's offer. An agreement is negotiated and recorded.

6.4.4 Actor or action and system response

Actor / Action	System Response
1. The data subject clicks a checkout link on the data collection entity's website, invoking the Agent Capability (in this case, the user's browser).	2. The Agent Capability requests the relevant conditional agreement object from the Negotiation Service (in this case, a CGI process on the data collection entity's Web server).
	3. The Negotiation Service obtains the conditional agreement object from the system (in this case, a CGI process accessing a back-end server).
	4. The system returns the conditional agreement object (in this case, an XML file).
	5. The Negotiation Service renders the terms of the conditional agreement object as an HTML form and returns the form to the Agent Capability.
	6. The Agent Capability presents the HTML form to the data subject.
7. The data subject enters the necessary PI, grants or denies the optional permission requests by clicking checkboxes, and approves the agreement by submitting the form.	8. The Agent Capability posts the form to the Negotiation Service (again, a CGI script on the web server).
	9. The Negotiation Service processes the form posted and submits the approved agreement (in this case, an XML file) to the system.
	10. The system stores the agreement object.
	11. The system acknowledges the Negotiation Service.

Continued…

Actor / Action	System Response
	12. The Negotiation Service sends a new HTML page to the Agent Capability.
	13. The Agent Capability (through the Interaction Service) presents the page to the data subject confirming a successful transaction and the negotiation is complete.

6.5 For further consideration

Agreement terms must be consistent with privacy or security policies. Privacy and security policies are necessary elements of an agreement, but they do not themselves present the user with specific choices about data collection and usage. To do that, a proposal or conditional agreement must present the actual agreement terms and permission choices available to the user, and these must be consistent with the data collection entity's privacy and security policies.

The Control Service must be able to store all the elements bound either into or by the agreement. The storage of the agreement object presents key requirements to the Control Service. If these elements are not stored in the agreement object itself, the Control Service must maintain the integrity of any external references to the actual PI data or permissions, wherever they are stored by the data collection entity. In addition, the Control Service must store the authentication credentials and modifiability attributes of permissions to support the requirements of the Access Capability.

Logging of a negotiation can itself produce PI. Even when a data subject or data subject agent wishes to remain anonymous or a negotiation fails to reach agreement; the negotiation process itself can produce PI that can potentially be associated with that data subject. Regulations or privacy policies must deal with this form of inferred PI.

7 Validation Service

7.1 Overview

The Validation Service is aimed at the partial satisfaction of the fair information practices principle of data quality. Personal information should be accurate, complete, and timely. Relevance is also suggested under data quality.

The aim of the Validation Service is to ensure data correctness at the time of entry, when the key issues are reliability, validity, and authenticity of the data. Also, data that has been stored for a length of time should be re-validated.

7.2 Functional description

The key functions of the Validation Service include the following:

- Provide consistency checks with corroboration. The consistency checks include bounds on parameter values and heuristic checks to which the data is subjected. Consistency checks can catch many input errors. The purpose of the corroboration feature would be to attempt to validate the data from independent sources. For example, a web search can access search engines to locate and access independent sources for some of the data, such as address and telephone number. Comparison checks could also be made against what is already in the database and what the user has input previously.
- Check for consistency against defined bounds and heuristics.
- Compare data being entered to both information that the subject has previously entered, and to related and supporting elements in the database.
- Access search engines to locate and access independent sources of some of the data, such as addresses and telephone numbers.
- Pass the data back to the Validation Service before being released to the requester if a significant amount of time has passed since the requested data has been validated.
- Verify the PI through to the Validation Service before it is added to the database if the PI originates from other than the data subject directly, as might be the case with derived PI or aggregated PI. The subject should also be able to view and correct their PI.
- Alert the system whenever the data fails the validity checks. This is particularly true when inconsistency or the nature of the failure suggests that the failure is not the result of error, but might be deliberate.

7.3 Actors and objects

The key actors and objects involved in the Validation Service are:

- **Data subject**, which provides personal and other information that requires validation.
- **Data objects**, which are personal and other information.

- **External corroboration entities**, which are used to check data entered by the subject.

7.4 Use case scenario

7.4.1 Name

External validation of data

7.4.2 List of scenario actors and objects

Data subject, external corroboration entity.

7.4.3 Scenario purpose and overview

Purpose: To validate personal data by accessing an external source.

Overview: The profile of personal information entered by the data subject is sent to the external identity authentication corporation called Equifax for the purpose of checking accuracy.

7.4.4 Actor or action and system response

Actor / Action	System Response
1. The data subject decides to provide personal information.	2. A data template is provided to the data subject.
3. The data subject completes the personal information template.	
	4. The template data is submitted to Equifax.
	5. The data is received; correctness of the data is confirmed.
	6. Notice is sent to the data subject confirming data correctness.
7. The data subject receives the confirmation.	
	8. The data is entered into a secure repository.
	9. An audit record is entered.

7.5 For further consideration

Accessing other resources: How can the Validation Service transparently access other sources for the validation and corroboration of PI?

Current credentials: How can credentials be kept current?

8 Access Capability

8.1 Overview

The Access Capability is motivated by the fair information practices principle of individual access. Data subjects should have access to any PI that a data controller or processor has collected about them, and have the ability to update or delete that PI as appropriate.

Access is not considered a formal service within the ISTPA Privacy Framework, but rather it is a capability that works with the Framework services, other system components, and security infrastructures. As an example, the Audit service may log each request to the Access Capability or to the Control Service for tracking. If a data subject believes there has been a breach of the agreement, the Access Capability may have the ability to invoke the Enforcement Service.

If an agent represents a data subject, the agent can provide the Access capability as part of its functions. The agent serves as the intermediary between the Interaction Service and the Control Service. The agent can be particularly useful to the data subject if it can also securely store the users' authentication credentials and interaction histories, as these make navigation and authentication easier for the user. If an agent represents the data controller or processor, this agent can serve as the Access Capability's interface to the Control Service.

The Access Capability can use the Negotiation Service (outbound from the data subject to the data requestor or third party) to establish and process communications with the data controller or data processor.

8.2 Functional description

The Access Capability provides a means for data subjects to view and modify the PI managed by a data controller or processor. The key functions of the Access Capability include the following:

- Provide a means for the data subject to locate the access mechanism provided by the data controller or processor (if necessary, using the Negotiation Service)
- Provide a means to identify and authenticate the data subject or PI owner.
- Provide a means to view the data subject's PI, including the agreement(s) negotiated with the data controller or processor.
- Modify or delete PI objects, preferences, or agreement as necessary.
- Confirm that modifications or deletions have been accepted, executed and recorded by the data controller, processor, certification authority or auditor.
- Provide access to the Enforcement Service or its Recourse function, if the data subject believes the terms of a privacy agreement have been violated.

8.3 Actors and objects

The key actors and objects involved in the Access Capability are:

- **Agreement object**, which consists of the completed agreement template after a successful negotiation.
- **Audit object**, which consists of PI processing events.
- **Auditor** or **investigative agent**, who audits or investigates data controllers or processors.
- **Certification authority**, which issues, suspends, and revokes certificates.
- **Control service data repository**, where the resulting agreement object is stored.
- **PI controller** or **processor**, which is responsible for the management and processing of PI.
- **PI object**, which consists of the PI authored by the data subject or PI owner.
- **PI owner** or **data subject**, who require access to audit services.
- **Preference object**, which details the data subject's privacy preferences governing the PI.
- **Privacy policy**, which details the data controller or processor privacy policy.

8.4 Use case scenario

8.4.1 Name

Individual access to operator-managed PI

8.4.2 List of scenario actors and objects

Agreement object, audit object, audit policy manager, control service data repository, data subject, PI controller, PI object, privacy policy, privacy preference object.

8.4.3 Scenario purpose and overview

Purpose: Data subject views and updates PI.

Overview: An individual wishes to view and update the PI currently stored in his or her account records that are managed by an e-commerce site operator, after the individual has made a purchase at the site.

8.4.4 Actor or action and system response

Actor / Action	System Response
1. The data subject invokes the Interaction Service (in this case, the user's browser) to request access to the data controller's website.	

Continued…

Actor / Action	System Response
2. The data subject navigates to the vendor’s Access Capability (in this case the “My Account” page of the website which accesses a CGI process).	
3. The data subject fills in an HTML form with a username and password for authentication.	4. The Access Capability submits the HTML-posted data to the Control Service (in this case, a CGI process accessing a back-end server), if necessary, using the Negotiation Service
	5. The Control Service returns authorization and a menu of agreement objects
	6. The Access Capability renders these as a web page to the Interaction Service.
7. The data subject navigates the menu to select an agreement object.	8. The Interaction Service requests this agreement object from the Access Capability.
	9. The Access Capability requests the agreement object from the Control Service.
	10. The Control Service returns the Agreement object to the Access Capability.
	11. The Access Capability renders the agreement object as an HTML form, with the relevant editing options, to the Interaction Service.
12. The data subject views and edits the PI object and agreements as desired, then submits the HTML form.	13. The Interaction Service posts the form to the Access Capability.
	14. The Access Capability processes the form data and passes it to the Control Service.

Continued…

Actor / Action	System Response
	15. The Control Service applies the necessary rules processing to validate that the revised or terminated agreement object is valid and acceptable to the vendor. (If not, the relevant steps are repeated, with feedback to the user about the error condition). Note that modification or deletion of certain PI or agreements may require re-negotiation of the agreement according to the terms of the original agreement. In this case, the Negotiation Service is called.
	16. The Control Service saves (or deletes) the appropriate agreement object and returns an acknowledgment to the Access Capability.
	17. The Access Capability returns a web page to the Interaction Service with confirmation to the data subject, completing an individual access cycle.
	18. The Audit Service, which provides logging, records the actions of the Access Capability or Control Service, as necessary.

8.5 For further consideration

Strong authentication poses a key usability challenge. Users already have a hard time remembering usernames and passwords. Remembering and managing stronger authentication credentials presents an even tougher problem. Agents are one potential solution to this problem.

Access is not limited to PI, but to the authentication data and permissions associated with that PI in an agreement. Access and modifications to authentication data or agreements may have higher security or business process impacts than the PI itself.

Modification or deletion of data or agreements by the user can have complex ramifications. For example, certain PI or agreements may be associated with services provided by a vendor (for example, a store discount card) whose terms state that the vendor has use of certain data. Deleting this PI or agreements could result in those services being cancelled. This can have even more complex ripple effects when these dependencies span multiple agreements. Informing users about these consequences and allowing them to make informed choices becomes a key requirement.

Unique keys or identifiers can simplify updating of PI but create privacy concerns and challenges. If a data subject has multiple agreements with a data collection entity, the use of a unique key or identifier in an agreement to match an instance of a PI element or elements can make updating much simpler and less error-prone for both the data subject and data controller or processor. However, these unique identifiers can be abused and pose additional threats to citizen privacy and informational self-determination.

9 Agent Capability

9.1 Overview

The Agent Capability is a software process that acts on behalf of a data subject or a requestor in order to support one or more of the services and capabilities defined in this Framework. Agent also refers to the data subject in the case of a manual process. The fair information practices that motivate the Agent Capability are notice and awareness, choice and consent, collection limitation, and openness.

The Platform for Privacy Preferences (P3P)-based agent, as defined by the World Wide Web consortium (W3C), will be used as an extended example below. However, the functionality for the Agent Capability in the ISTPA Privacy Framework is more broadly defined. For example, the Agent Capability supports the Negotiation Service and interacts with the Control Service for privacy policy decisions, such as allowing use of the data subject's email address for a newsletter.

Although the P3P-based agent is used for convenience in the scenario below, it does not necessarily indicate a formal endorsement of the W3C's P3P by the ISTPA. Part of the ISTPA charter stipulates examining the usability, transparency, and extensibility of existing and future privacy technologies, including P3P.

P3P simply serves as an example of the requirement placed on any computer-based privacy management system to define a grammar for describing data subject preferences and subsequent agreements and permissions in machine-readable form.

9.1.1 What is an intelligent agent?

Broadly speaking, an agent is a software entity that is persistent, can perceive, reason about, and affect its environment, and it can communicate with other agents and humans. Most agents have some degree of autonomy. Some agents are mobile; that is, they can move across networks to execute on other computers. Agents can collaborate to solve problems with other agents or with humans, allowing a group of task-specific agents to solve complex problems.

The type of agents that this Framework is most interested in, however, are intelligent; they have explicit knowledge about the tasks that they will perform. Intelligent agents are rapidly gaining popularity in this increasingly networked world because they can provide information monitoring, searching, filtering, and decision support. This support and filtering lets humans focus more clearly on the information that is most relevant or important to them in a given context. Intelligent agents can also be personalized so that they meet specific needs that users have. Agents can monitor complex situations, alerting users to necessary actions only when needed.

More specifically, privacy agents are intelligent agents that have the knowledge and authority to help users make decisions about how, when, and why they want to or need to share personal data.

9.1.2 Website agents for negotiating privacy policies

A given privacy policy might allow a website to collect the minimum amount of information necessary to do business, to not keep it longer than necessary, and to not use it for other purposes. This is a simple privacy exchange. The website sends the privacy policy, and if users do not agree to it, the website does not do business with them.

More interesting, however, is the possibility of negotiating a privacy and data exchange agreement that allows for the personalization of relationships between customers and business. Website agents might have a complex set of rules for ensuring that the website gets the correct, "tailored" information from each user, that it can provide personalized shopping to the user, and that it "pays" only what it wishes for any piece of information. This agent must be able to integrate personal information across sources (including privacy agreements), know what information it has already collected, and what information it may not want to pay for again unless it has evidence that the information may have changed. It will only be possible to tailor a relationship with a user who is pre-identified, perhaps from a continuing relationship such as a loyalty card.

This suggests the possibility for an initial agreement that will include a small data exchange for the purpose of allowing some tailoring of the negotiation. For example, the initial announcement of a privacy agent at the website can give some jurisdictional information so that the privacy negotiation will be consistent with local, legal, or cultural constraints.

9.1.3 Accommodating businesses that operate beyond the web

Protecting personal privacy on the Internet is one issue, but protecting personal privacy at every customer touch point is an issue with which enterprises are concerned. Businesses want a consistent relationship with the customer, and customers expect consistent privacy and data exchange policies and behaviors throughout the enterprise. The following scenarios present ways to extend agents into spaces beyond the web.

- A kiosk with a smart card reader; for example, a frequent shopper card. Users' privacy preferences are stored on the card, or the card is used to retrieve the preferences from a repository that is maintained by the business or by an intermediary. Data is collected from the transactions and interactions with the kiosk.
- A terminal, where cashiers swipe the smart card; essentially the same as the kiosk example.
- An infomediary approach, where the personal data repository is maintained by the infomediary that also houses consumers' personal privacy agents.
- A frequent shopper program, where users pre-specify their preferences. Retailers interact with the agent that the users can use at any time to update their privacy preferences through a wireless mobile device.
- Data that has been collected from all touch points and interactions, along with related privacy agreements, is integrated in a data warehouse. The enterprise can then use the data warehouse to administer its adherence to privacy agreements and privacy policies, including integrating data across all customer touch points and interactions.

9.2 Functional description

Consider the use of personal privacy agents in a typical use case scenario where consumers delegate their privacy and data exchange decisions and negotiations to a personal privacy agent. The consumers need to first receive the basic specifications from a trusted source, and be able to customize the rules to meet their own needs and preferences. At that point, retailers can tailor negotiated data exchanges to personalize interactions with the consumers, and request specific types of information from special categories of consumers. These exchanges can help address retailers' most pressing business questions: although the retailer may have to pay for the data with discounts or frequent shopper points, there is a much higher likelihood that the collected, specific data is actionable.

Users might want to have access to a set of flexible rules that would let them customize how much personal information is revealed. For example, three possible rules might include one where users allow the exchange of only small amounts of personal information, a second which allows for e-commerce and personalized relationships with a set of trusted sites, and a third for unlimited exchange or unrestricted internet interactions.

These rules could be represented in a formal grammar by the providing organizations. Users' agent reasoning engines would interpret the rules and compare them with privacy proposals from a given website. The configuration interface would allow users to easily make changes to the set of rules that they have chosen, and to tailor them to their own needs. It could also represent a set of trusted websites and businesses, so that users can have a different set of rules for business with which they have an ongoing relationship.

The key functions of the Agent Capability include the following:

- Provide persistent storage for PI objects, preferences, and agreements.
- Provide a means for the data subject to specify and enter PI objects, preferences, and rules governing PI object access and processing.
- Have the ability to act on behalf of the data subject in conducting rule-based negotiations with data controllers and processors.
- Provide a means to interact with Audit, Certification and Security Services to manage and monitor stored PI objects, preferences and agreements.
- Process and securely execute rules governing PI objects, preferences and agreements.
- Provide a means to communicate and/or be able to execute a protocol(s) for interacting with other agents or services.
- Provide a means to transport PI objects, preferences and agreements or perform mobile operations across networks.

9.3 Actors and objects

The key actors and objects involved in the Agent Capability are:

- **Agent object**, which contains PI objects, preferences and agreements.
- **Agreement object**, which consists of the completed agreement template after a successful negotiation.
- **Audit object**, which consists of PI processing events.
- **Certification authority**, which issues, suspends, and revokes certificates.
- **Control service data repository**, where the resulting agreement object is stored.
- **Meta dictionary**, which defines and structures PI elements, audit events, rules and nomenclature.
- **PI controller** or **processor**, which require access to agent capabilities.
- **PI owner** or **data subject**, which require access to agent capabilities.
- **PI object**, which consists of the PI authored by the data subject or PI owner.
- **Preference object**, which details the data subject's privacy preferences governing the PI.
- **Process certificate object**, which binds the necessary credentials to the data controller or processor and points to issuing authorities examination report.

- **Process certificate manager**, which ascertains relevant privacy principles and/or regulatory requirements, assembles the necessary criteria for process certification, and manages the certificate life cycle.
- **Privacy policy**, which details the data controller or processor privacy policy.

9.4 Use case scenario

9.4.1 Name

Multi-round agreement negotiation

9.4.2 List of scenario actors and objects

Agreement, data subject, data subject agent, PI object, privacy policy, privacy preference object, website agent.

9.4.3 Scenario purpose and overview

Purpose: Data subject agent negotiates an agreement with website agent.

Overview: The Interaction Service will be needed for the agents' interaction with the user, with a website, with an enterprise's agent, and with the humans at the enterprise. The Negotiation Service will facilitate the negotiation of privacy agreements, where applicable. The Certification Service may be needed to ensure the identity of each participant in a transaction.

9.4.4 Actor or action and system response

Consider the following exchange, which illustrates using privacy policies to enhance customer-business relationships and personalization while protecting personal privacy. This illustration makes use of multi-round negotiation. The exchange only takes a few seconds and would be completely invisible to the consumer.

Actor / Action	System Response
1. The data subject goes to a website to make a gift purchase.	
2. The data subjects' personal privacy agent announces itself at the website.	3. The website agent accepts the data subject agent for privacy and data exchange negotiation.

Continued…

Actor / Action	System Response
4. The data subject agent asks for a privacy proposal.	5. The website agent sends a proposal requesting physical contact information and personal preference in exchange for 10% off on the next purchase. The website agent is configured to pay for preference information from users who are shopping on a holiday.
6. The personal privacy agent is configured to never expose or send preference data. Instead, the agent sends a counter-proposal agreeing to send physical contact information in exchange for 5% off on the next purchase, since the user has categorized the site as a trusted site.	7. The website agent accepts the counter proposal.
8. The user goes to the website and shops, after sending the negotiated information.	

9.5 For further consideration

How will consumers react to and interact with a privacy agent?

How much detail about data sharing will users want to understand? How feasible would it be to create an adaptive privacy agent that would provide varying levels of interaction based on how much interest different users have in customizing their settings?

How can agents save people time and allow them to have confidence in the decisions being made?

What kind of interface is most appropriate for a user to configure the agent, and by what means should the agent communicate with the user?

10 Usage Capability

10.1 Overview

The fair information practice governing use limitation motivates the Usage Capability. This service assumes the role of “processing monitor.” It ensures that the active use of PI that is outside the direct control of the subject complies with the terms and policies of any agreement and applicable regulations. Such usage models include transfer, derivation, aggregation, pseudo-anonymization, linking, integration, and inference.

10.2 Functional description

The key functions of the Usage Capability include the following:

- Provide the guidelines, controls, and allowed processing for all uses of personal information (PI) as well as potentially identifiable personal information, even when the data is outside direct control of the subject. Well-known versions of the Usage Capability include data mining, profiling individuals, market research, and contact or sales list processing of various types. The Framework is responsible to either provide or hide personal information based on the user agreements, laws, and policies.

Other requirements for the Usage Capability vary, depending on the data models being utilized. These data usage models include the following defined categories and their privacy usage requirements:

- **Transferred:** PI is requested and subsequently transferred to the data requestor, who is faithful to the permissions of the subject.

  Example: The data subject browses to a web service and signs up for a subscription service. As part of the transaction, PI is provided to the web service under permissions spelled out in the privacy policy of the service. Subsequent use of the PI is consistent with the permissions. Any exceptions are reported to proper authorities.

- **Aggregation / Depersonalization:** All personal information is deleted from the data, allowing the data to be analyzed in the aggregate through data mining processes. Data is freely usable without restriction.

  Example: After the deletion of direct personal information, such as name and address, credit card transaction histories of cardholders, along with demographic information, can be analyzed to determine buying patterns, preferences, and other information valuable to merchants and advertisers.

- **Pseudo-anonymization / De-identification:** Personal information is replaced by a non-identifiable linkage record in order to prevent the using entity from being able to identify the individual. A trusted third party maintains the information needed to connect the linkage record back to the individual under the controls of the existing user agreements, laws, and policies applicable to that person’s information.

  Example: Health records used for medical research can have personal information replaced by an anonymous linkage record to protect the identity of the individuals. If, during the usage of the associated data, a valid requirement arises to be able to determine the identity of the individual, the linking record and identity can be supplied to a third party who is responsible for validating the requirement for identification based on the existing laws, policies and preferences.

- **Derivation:** Personal information and associated data about an individual may be used by entities to create or 'derive' new information about the individual. This derived data is generally added to the secure repository and should be subject to the existing user agreements, laws and policies.

  Example: Banks, credit bureaus and financial institutions create credit ratings about individuals based on past dealings, such as payment histories, income levels, length of time in a job, and other data. This credit rating information becomes part of the personal information in the bank's data repository and is frequently shared between or submitted to other credit rating entities.

- **Extension / Linkage**: In this usage model, new data about an individual from a range of sources is added to or extends the existing personal information. Also, multiple records of personal information can be linked together to provide a broader set of data about an individual.

  Example: Two companies merge or form an alliance in order to provide a more comprehensive set of services to their combined user base through linking their customer data repositories and files. For instance, a bank and an insurance company merge and then link the information about their customers so that both the financial and insurance profiles are available to do marketing, rating, and other data usage activities.

- **Inference:** By collecting data from multiple sources, some of which provide personal information, it may be possible to analyze the data in ways to infer the identity of individuals represented in the non-PI data and then link the various sources of data to that specific individual.

  Example: Users may be tracked on the Internet using cookies to identify them and the locations they visit. It is possible to have one reference, which includes personal information to be combined with other sources of information, which use only cookies to determine the identity of the individual and subsequently link all the data together. Any user agreement information must apply to all of the linked data about the inferred individual. Lacking any user agreement data, the user should be contacted to obtain an agreement before usage of the personal information is allowed. In addition, all laws and policies also apply to any linked data.

- **Integration:** In this model, the base records of personal information are regularly updated with new and additional information obtained from various sources. The data is integrated into the primary repository files and records for each individual.

  Example: Credit bureaus integrate data from many sources into a single record or credit history for individuals. These integrated records are subject to all pertinent user agreements, laws, and policies.

Throughout the life cycle of personal information, the permissions granted by the subject of the PI must always apply or be re-negotiated. For that purpose, the permissions must be logically bound to the PI.

10.3 Actors and objects

The key actors and objects involved in the Usage Capability are:

- **Data subject.**
- **Data collectors/controllers.**
- **Personal Information** (includes any transferred, derived, aggregated, anonymized, linked, integrated, and inferred data).

10.4 Use case scenario

10.4.1 Name

Applying permissions to integrated PI.

10.4.2 List of scenario actors and objects

Data subject, data collectors/controllers

10.4.3 Scenario purpose and overview

Purpose: Apply the original subject permissions to integrated PI.

Overview: On two different occasions, PI is requested and transferred to a data requestor from a data subject. The PI is integrated, but the separate permissions are used to govern use of the separate PI.

10.4.4 Actor or action and system response

Actor / Action	System Response
1. The data subject interacts with a website for services and exchange of PI with associated permissions.	2. The system transfers the PI and permissions to the data requestor.
3. Subsequently, the data subject interacts with the same website and grants permissions to a different set of PI.	
	4. Both PI sets are logically integrated within the website's privacy framework, but the permissions are kept separated.

Continued…

Actor / Action	System Response
	5. The website uses the integrated PI, but consistently with the permissions and subject to the jurisdictional constraints of the site.
	6. The audit log records the activities.

10.5 For further consideration

Maintaining personal privacy on the Internet: Data usage about individuals, entities, and groups is at the root of the worldwide debate about privacy. None of the issues are new; what is new is the extraordinary capability brought about by the Internet to collect, exchange, and use information in ways which are beyond the existing laws, policies, and agreements. It is clear that new user agreements, laws, and policies must be put in place to recognize and deal with these issues.

Agreement enforcement: Data from repositories for usage applications is frequently extracted and then separated from direct control of the subject. How the existing user agreements, laws, and policies continue to be enforced in these circumstances is a major issue. Should user agreements always be linked with user data, even when personally identifiable information has been deleted?

PI linking through mergers: Mergers of major companies such as banks and insurance companies, in some cases for the primary reason of combining customer bases, create the opportunity to extend and link far more information about individuals than was previously possible. Should there be laws, policies, and new user agreements put in place to govern these new linkages and uses of the data about individuals?

Derived data: Does derived data about an individual have to be made available for review by that individual or is the information owned and proprietary to the entity that derived the data?

Individual access to PI: Do users have a right to understand what uses have been made of their information and the entities who have accessed their data? Today, a user can obtain a record of who has requested their credit history. Should that be the same model for all personal information?

Maintaining business PI: While the primary focus of the Privacy Framework is individuals, these principles also apply to other entities such as companies, groups, and other identifiable organizations that may have a right to privacy under the laws and policies. An example was the challenges by corporations who were 'profiled' by Amazon.com on their employee's book buying preferences.

Glossary

Access control: Restricts access to data or services to a particular identity or group of identities. Access control can be either discretionary or mandatory. Access control lists (ACLs) are typically used for discretionary controls. Labels, which indicate the subject's clearance, are used for mandatory controls.

Actor: Any entity that send or requests information to or from a service. This may be an individual or machine, or a corporate or government entity.

Administrator: An entity that can override the access rights for a system. It can change settings and grant others all or a subset of access rights to the system. An administrator, however, is not the same as an auditor.

Aggregation (depersonalization): All personal information is deleted from the data, allowing the data to be analyzed in the aggregate through a data mining processes.

Agreement object: A collection of PI combined with set of agreements that can either restrict or allow usage of the PI.

Anonymous PI: The dissociation of PI across a sufficiently large population such that no PI can be associated with a particular data subject.

Audit: A chronological record of events. Audits are typically used to provide accountability.

Auditor: The entity that sets audit controls (i.e., what is being audited) and has access to the audit records or logs. The auditor is not the administrator; the administrator is often the subject of the audit.

Authentication: Verification of the evidence or proof of a claimed identity.

Authorization: The mechanism that a system uses to provide access control over data or services.

Certificate: A sequence of data providing identity or attributes for an entity. It is usually signed by a trusted entity. An example is an x.509 certificate.

Credential: The representation of an authentication.

Data collection entity: An entity that either requests PI directly from the data subject or collects agreement objects from other data collection entities.

Data object: The actual data, whether PI or other, which is passed into or out of a service.

Data subject: The individual from whom information is gathered or to whom information is directly associated. This is also the "PI owner."

De-identification (pseudo-anonymization): Personal information is replaced by a non-identifiable linkage record in order to prevent the using entity from being able to identify the individual. A trusted third party maintains the information needed to connect the linkage record back to the individual under the controls of the existing user agreements, laws, and policies applicable to that person's information.

Depersonalization (aggregation): All personal information is deleted from the data, allowing the data to be analyzed in the aggregate through a data mining processes.

Identification: A claim identity by an entity. This usually involved associating a unique label to an entity or set of entities such as a person, machine, process, or application. There does not need to be a direct association between the entity and the unique label. An alias may be used.

Judicial authority: A set of rules established by a governmental body that has jurisdiction over the subject.

PI (personal information): Information that is directly associated with a data subject, such as name, email or physical addresses, government identification numbers, health information, etc.

PI controller: Any entity that holds or controls PI.

PII (personally identifiable information): Information that associates a particular PI or set of PI to a data subject.

Preference object: An object that details the data subject's privacy preferences that govern how the PI is used.

Privacy: The proper handling and use of personal information throughout its life cycle, consistent with data protection principles and the preferences of the subject.

Pseudo-anonymization (de-identification): Personal information is replaced by a non-identifiable linkage record in order to prevent the using entity from being able to identify the individual. A trusted third party maintains the information needed to connect the linkage record back to the individual under the controls of the existing user agreements, laws, and policies applicable to that person's information.

Pseudonym: An identity that is an alias of a data subject.

Regulatory authority: Usually a non-government set of rules established by a trade organization or other associations.

Service: A functional unit defined by the ISTPA Privacy Framework that performs actions on data objects (e.g., PI, agreement objects) either by the direct request of an actor, or by a process.

Actors and Objects

- AUDIT:

- **Agreement object**, which consists of the completed agreement template after a successful negotiation.
- **Audit object**, which consists of PI processing events.
- **Audit policy manager**, which matches appropriate privacy principles, practices and regulatory requirements that are necessary for PI-protected and compliant processing.
- **Auditor** or **investigative agent**, which require access to audit services.
- **Certification authority**, which issues, suspends, and revokes certificates.
- **Control Service data repository**, where the resulting agreement object is stored.
- **Meta dictionary**, which defines and structures PI elements, audit events and nomenclature.
- **PI controller** or **processor**, which require access to audit services.
- **PI object**, consisting of PI that was either authored by the data subject or the PI owner.
- **PI owner** or **data subject**, which require access to audit services.
- **Preference object**, which details the data subject's privacy preferences governing the PI.
- **Privacy policy,** which details the data controller's or processor's privacy policy.
- **Regulatory** or **judicial authority**, which require access to audit services.
- **Third parties**, who may wish to process PI or to gain unauthorized access or processing privileges by assuming the identity of legal PI agents or processors.

- CERTIFICATION:

- **Agreement object**, which consists of the completed agreement template after a successful negotiation.
- **Audit object**, which consists of PI processing events.
- **Audit policy manager**, which matches appropriate privacy principles, practices, and regulatory requirements necessary for PI protection and for compliant processing.
- **Auditor** or **investigative agent**, which requires access to the audit services.
- **Certification authority**, which issues, suspends, and revokes certificates.
- **Control Service data repository**, where the resulting agreement object is stored.
- **Meta dictionary**, which defines and structures PI elements, audit events and nomenclature.
- **Personal Information (PI) controller** or **processor**, which requires access to the audit services.
- **PI object**, which consists of the PI that is authored either by the data subject or by the entity that owns the PI.
- **PI owner** or **data subject**, who requires access to the audit services.

- **Preference object**, which details the data subject's privacy preferences governing the PI.
- **Process certificate object**, which binds the necessary credentials to the data controller or processor and points to issuing authorities' examination report.
- **Process certificate manager**, which ascertains relevant privacy principles and/or regulatory requirements, assembles the necessary criteria for process certification, and manages the certificate life cycle.
- **Privacy policy**, which details the data controller or processor privacy policy.
- **Regulatory authority** or **judicial authority**, which requires access to the audit services.
- **Third parties**, who may either wish to process the PI legally or possibly gain unauthorized access or processing privileges by assuming identity of legal PI agents or processors.

■ CONTROL:

- **External requestors.**
- **Internal requestors.**
- **Personal information (PI).**
- **Privacy policies.**

■ ENFORCEMENT:

- **Agreement object**, which consists of the completed agreement template after a successful negotiation.
- **Audit object**, which consists of PI processing events.
- **Audit policy manager,** which matches the appropriate privacy principles, practices, and regulatory requirements that are necessary for PI protection and for compliant processing.
- **Certification authority,** which issues, suspends, and revokes certificates.
- **Control service data repository**, where the resulting agreement object is stored.
- **Personal Information (PI) controller** or **processor**, which requires access to the audit services.
- **PI object**, which consists of the PI that is authored either by the data subject or by the entity that owns the PI.
- **PI owner** or **data subject**, who require access to the audit services.
- **Preference object**, which details the data subject's privacy preferences governing the PI.
- **Privacy policy,** detailing the data controller's or processor's privacy policy.
- **Process certificate object**, which binds the necessary credentials to the data controller or processor, and points to the issuing authorities' examination report.
- **Process certificate manager**, which ascertains the relevant privacy principles and/or regulatory requirements, assembles the necessary criteria for process certification, and manages the certificate life cycle.
- **Regulatory authority** or **judicial authority**, which requires access to the audit services.

- **Third parties,** who may either wish to process the PI legally or possibly gain unauthorized access or processing privileges by assuming identity of legal PI agents or processors.

- ■ INTERACTION:

- **Data collection entity**, which acts as a requestor of data.
- **Data subject**, which is the source of the PI.
- **Extra-Framework entities**, which function as correspondents with the Framework.
- **Data objects**, which are any data that can be passed into and out of the Framework.

- ■ NEGOTIATION:

- **Agreement duration**, which is the period of time the agreement will remain in effect.
- **Agreement object**, which results from a successful data collection negotiation.
- **Agreement PI object requests**, optional or required, which specify the PI elements that have been requested by a proposal.
- **Agent**, which represents respectively the data subject and the data collection entity entering the negotiation.
- **Authentication credentials**, which are used to validate the identity of the parties.
- **Conditional agreement object**, which represents a proposal object whose terms at least one party has agreed to, subject only to approval of the same terms by the other party. Note that a proposal object that the data collection entity will accept a valid response for is already a conditional agreement object.
- **Control service data repository**, where the resulting agreement object is stored.
- **Data collection purpose(s),** which are the stated reason(s) for the collection of PI from a data subject.
- **Optional permission requests**, optional or required, which are requests to use the PI that is covered by an agreement for a purpose other than the primary purpose(s) defined in the purpose element above. A permission contains the same elements as an agreement except policies and authentication credentials. These are defined by the containing agreement. It can be considered "an agreement within an agreement."
- **PI controller** or **processor,** which controls or processes PI.
- **PI object**, which consists of the PI authored by the data subject or PI owner.
- **PI owner** or **data subject**, the individual or entity who owns and authors the PI.
- **Policy references**, which are used to assert the privacy, security, and other relevant policies of the data collection entity covering this agreement.
- **Proposal object**, which contains authentication credentials, policy preferences, data use purposes, agreement duration, requested PI objects and permissions.
- **Proposal object repository**, which stores proposal or conditional agreement objects on behalf of the data collection entity. (Note that this may be managed by the Control Service.)

- VALIDATION:

- **Data subject**, which provides personal and other information that requires validation.
- **Data objects**, which are personal and other information.
- **External corroboration entities**, which are used to check data entered by the subject.

- ACCESS:

- **Agreement object**, which consists of the completed agreement template after a successful negotiation.
- **Audit object**, which consists of PI processing events.
- **Auditor** or **investigative agent**, who audits or investigates data controllers or processors.
- **Certification authority**, which issues, suspends, and revokes certificates.
- **Control service data repository**, where the resulting agreement object is stored.
- **PI controller** or **processor**, which is responsible for the management and processing of PI.
- **PI object**, which consists of the PI authored by the data subject or PI owner.
- **PI owner** or **data subject**, who require access to audit services.
- **Preference object**, which details the data subject's privacy preferences governing the PI.
- **Privacy policy**, which details the data controller's or processor's privacy policy.

- AGENT:

- **Agent object**, which contains PI objects, preferences and agreements.
- **Agreement object**, which consists of the completed agreement template after a successful negotiation.
- **Audit object**, which consists of PI processing events.
- **Certification authority**, which issues, suspends, and revokes certificates.
- **Control service data repository**, where the resulting agreement object is stored.
- **Meta dictionary**, which defines and structures PI elements, audit events, rules and nomenclature.
- **PI controller** or **processor**, which require access to agent capabilities.
- **PI owner** or **data subject**, which require access to agent capabilities.
- **PI object**, which consists of the PI authored by the data subject or PI owner.
- **Preference object**, which details the data subject's privacy preferences governing the PI.
- **Process certificate object**, which binds the necessary credentials to the data controller or processor and points to issuing authorities examination report.
- **Process certificate manager**, which ascertains relevant privacy principles and/or regulatory requirements, assembles the necessary criteria for process certification, and manages the certificate life cycle.
- **Privacy policy**, which details the data controller's or processor's privacy policy.

- USAGE:

- **Data subject.**
- **Data collectors/controllers.**
- **Personal Information** (includes any transferred, derived, aggregated, anonymized, linked, integrated, and inferred data).

Useful Links

In addition to the specific words introduced in this document, definitions for standard industry terms can be found at the following websites:

- ISO Subcommittee SC27 (security) glossary:

  http://www.din.de/ni/sc27/doc6.html/

- National Computer Security Center – security glossary:

  http://www.fas.org/irp/nsa/rainbow/tg004.htm

- Center for Democracy and Technology (CDT) Guide to Online Privacy – Glossary:

  http://www.cdt.org/privacy/guide/terms/

- NSA Glossary for Security and Intrusion Detection:

  http://www.sans.org/newlook/resources/glossary.htm

- IETF RFC 2828 – Internet Security Glossary:

  http://www.faqs.org/rfcs/rfc2828.html

- TRUSTe Privacy Glossary:

  http://www.truste.org/partners/users_glossary.html

Unified Modeling Language (UML) and related websites (partial list):

- OMG Unified Modeling Language (UML) Specification:

  http://www.omg.org/technology/documents/formal/uml.htm/

- UML 1.4 with Action Semantics:

  http://www.omg.org/cgi-bin/doc?ptc/2002-01-09/

- Kendall Scott's UML Dictionary:

  http://usecasedriven.com/UML.htm

- OMG UML Resource Page:

  http://www.omg.org/technology/uml/index.htm

- Cetus Links on Objects & Components (UML):

  http://www.cetus-links.org/oo_uml.html

- Cris Kobryn's UML Forum site:

  http://www.celigent.com/uml/

- Rational UML Resource Center:

  http://www.rational.com/uml/

Other useful sites (partial list):

- Kidz Privacy:

  http://www.ftc.gov/bcp/conline/edcams/kidzprivacy
- Electronic Privacy Information Center:

  http://www.epic.org
- The Office of the Federal Privacy Commissioner:

  http://www.privacy.gov.au
- Privacy Foundation:

  http://www.privacyfoundation.org
- Safe Harbor:

  http://www.exports.gov/safeharbor
- Privacy Commissioner of Canada:

  http://www.privcom.gc.ca
- Privacy International:

  http://www.privacyinternational.org
- Privacy Exchange:

  http://www.privacyexchange.org
- BBB On-Line:

  http://www.bbbonline.org/consumer
- CPA WebTrust:

  http://www.cpawebtrust.org
- Privacy.net - The Consumer Information Organization:

  http://www.privacy.net
- Privacy Council:

  http://www.privacycouncil.com/
- Americans for Computer Privacy:

  http://www.computerprivacy.org