

The New
Federated Privacy Impact Assessment (F-PIA):
Building Privacy and Trust-Enabled Federation

January 2009

The New Federated Privacy Impact Assessment (F-PIA): Building Privacy and Trust-Enabled Federation

Foreword

Throughout my career as a privacy professional and as the Information and Privacy Commissioner of Ontario, Canada, I have always advanced the view that, “privacy is good for business – good privacy is good business.” A lack of attention to privacy can have a number of adverse consequences for businesses, ranging from damage to reputation and brand to, most important, loss of customer trust and loyalty. This problem has become a significant one for organizations since consumers increasingly face the ever-growing threat of identity theft, as just one example of the potential misuses of information in online activities. This is due, in large part, to the fact that the Internet was not created to deal with individuals who intend to commit fraudulent or malicious actions. To ensure the continuation of a robust and trusted online technology ecosystem, we desperately need an ability to distinguish identity thieves from legitimate users: enter “Identity Management.”

I issued my first publication on Identity Management in 2006, on the *7 Laws of Identity: The Case for Privacy-Embedded Laws of Identity in the Digital Age*.¹ This work took the concept of an “identity layer” for the Internet (a broad conceptual framework for a universal, interoperable identity system) and focused on how privacy-enhancing features can and must be embedded into the design of a universal identity system architecture.

Since that time, the Internet has continued to evolve. “Web 2.0,” as today’s iteration of the Web is more commonly referred to, has the potential to provide convenient time-saving services, tailored to the individual user. The privacy concerns revolve around the demands for greater amounts of personal information in this “new Internet” for use by multiple parties, to authenticate a user’s identity; further, these exchanges of identity information may not always directly involve the individual. Privacy involves providing individuals with appropriate control over their personal information and ensuring that adequate safeguards are in place to protect the information.² This is where Federated Identity Management (FIM) comes

1 Available at: www.ipc.on.ca/images/Resources/up-7laws_whitepaper.pdf

2 There was a time when it was believed that the more an individual had granular control over information, the more he or she could control privacy. While this is true in theory, highly detailed and fine-grained controls over information tend to overwhelm the average user because of the complexity and frequency of choices available. Thus, the true solution may involve “appropriate” controls, as determined by the context.

into play. FIM allows consumers to simultaneously and securely sign on to the networks of more than one enterprise, for the purpose of conducting various transactions, while still maintaining their privacy.

For some time now, my staff have participated as invited experts in the Public Policy Expert Group (PPEG), the Project, working alongside other member organizations, to advance privacy-enabled technologies. Together with Joseph H. Alhadeff, Vice-President for Global Public Policy and Chief Privacy Officer at Oracle Corporation (who provided significant expertise and input on behalf of the Liberty Alliance), we saw the need for a resource to assist organizations in designing privacy into the process at an early stage, as they embarked on developing a federated system for identity management. I agreed to participate in this joint venture because I was attracted to Liberty Alliance's holistic approach toward including privacy in the development of identity standards. In addition, its reputation for assisting organizations in establishing identity ecosystems that operate responsibly and securely, using decentralized authentication (so that a user's personal information does not have to be centrally stored), was a real plus.

From a practical perspective, I believe that FIM holds the promise of serving as the fundamental element in creating an identity layer to which consumers can entrust their personal information. Privacy protection is not available in a standard one-size-fits-all model. Each business is unique, and privacy needs are equally unique, which is why I encourage businesses to develop a "culture of privacy." By a culture of privacy, I mean developing a "mindset" – a way of thinking throughout the organization that is committed to better information management practices that are respectful of privacy. Even the most advanced technologies, coupled with the most rigorous privacy policies, will not be effective if they do not become an accepted part of your business culture. This white paper will assist organizations to achieve greater user trust in such a federation – an admirable goal. At its essence, it is a practical assessment of how privacy can be applied to a group of organizations and businesses that wish to create a community to manage their clients' identity – one that is based on trust.

Ann Cavoukian, Ph.D.
Information and Privacy Commissioner of Ontario
Canada

1 Introduction

1.1 The Web 2.0 World

Our world is becoming increasingly interconnected. Distributed networks of service and information providers are operating across global information and value chains. Observing, with greater frequency, dense inter-networking, large-scale data sharing, and the constant evolution of relationships between organizations, it becomes clear that firms are moving from “multinational” to “global” in nature, and that the concept of enterprise has morphed into the concept of an ecosystem.³ We are seeing the emergence of more than just Web 2.0 – we are, in fact, seeking out the World of 2.0.

In the online realm, new Internet services are exploiting the potential to provide detailed personalized services to individuals in every facet of their information-laden lives.⁴ The benefits related to these services, however, create information trails and broader possibilities for unauthorized access that require new considerations related to information protection. In these new paradigms, personally identifiable information becomes digitized, routed and processed on high-speed, high-capacity networks that are independent of each other, with fewer traditional hallmarks of information collection and control. New trust models are now being considered to deal with these information flows across distinct organizations at the ecosystem level (see our recently published white paper regarding the privacy implications of digital identity).⁵ Concepts such as cloud computing, in which organizations share both data and processing resources in order to coordinate a business process, are also beginning to take hold, creating new opportunities for cross-ecosystem industry collaborations.

As a first step toward creating trust in this new inter-networked or “federated” model, organizations are building secure ways in which to collect and store personal information. There are a number of technical means to ensure that information moving between organizations is securely transferred. However, technology is only part of the solution. Technology operates in support of people, policies and procedures, and in conjunction with the legal instruments that bind parties to the obligations related to the appropriate deployment of technology. As these enterprises and organizations start developing ecosystem-based rules and procedures, new tools will be needed to evaluate and oversee the deployment of technology, the implementation of policies and procedures, and the operation of contracts.

3 For more, see the IPC publication *Privacy and the Open Networked Enterprise*, December 2006, at: <http://www.ipc.on.ca/images/Resources/up-opennetw.pdf>

4 See, for instance, *A View from 2018: A Glimpse of the Internet Future*, at: www.biac.org/members/iccp/mtg/2008-06-seoul-min/Final_View_from_2018_ICCP_Chair_Paper.pdf

5 See the IPC publication *Privacy in the Clouds*, May 2008, at: www.ipc.on.ca/images/Resources/privacyintheclouds.pdf

Such tools are necessary to demonstrate that systems that may be technically sound can also be trusted by individuals. This paper looks at a number of the concepts underlying FIM, important practice considerations, and one of the most important baseline tools needed for the trust-enabled ecosystem – the “Federated Privacy Impact Assessment,” or F-PIA.

1.2 Federated Identity Management (FIM)

As we move forward from an enterprise model, where individuals interact with companies they know, to an ecosystem model, where information is shared within and across enterprises and value chains for a variety of purposes, new ways of dealing with personal information related to identity must be created. Federated Identity Management (FIM) systems are emerging to fill this gap – to create the Internet’s missing identity layer.⁶

Within the FIM model, identity credentials issued to a user by a particular service or institution are recognized by a broad range of other services. Though complex to implement online, this is similar in concept to, and can provide improvements over, traditional identification schemes in the “physical world.” A typical example would be government-issued ID credentials (birth certificate, driver’s licence, passport, citizenship card, etc.) issued by an institution (a government agency) that is broadly recognized by others (as proof of name, address, age, etc.). The user of the service does not need to prove his/her identity with each transaction; rather, it is enough to show that he/she has, at some prior point, been authenticated by a trusted authority. The service’s burden then lies not in identification of the presenter but in the verification of presented credentials – a much less onerous task. The improvement in the online federated model is that systems can be architected to ensure that only the least amount of relevant data needed to establish a credential is provided to a requestor. Going back to the driver’s licence example, if the licence is being used for age verification, there is also other identifying information that may be seen or scanned, which is unrelated to age. Further, most age verification requirements relate to age ranges, and thus do not require review of an exact birthdate. In such a situation, an FIM-based validation can, unlike the licence, return only a statement that the credential’s holder is over 18, 21, or 65 –

6 For more information on the missing layer, see the IPC publication *7 Laws of Identity: The Case for Privacy-Enabled Laws of Identity for the Digital Age*, at: ipc.on.ca/images/Resources/up-7laws_whitepaper.pdf

whatever the relevant range – without revealing any additional information about the individual.⁷

Situated online, Federated Identity Management has the potential to allow individuals to use the same username and password, or other personal identification, to securely sign on to the networks of more than one enterprise, in order to conduct transactions while maintaining privacy protections (a process known as “single sign-on”). Federation can allow companies to share applications and information securely without the need to maintain full user accounts for their partners’ clients (a helpful privacy best practice). Open technology standards enabling FIM can also ensure choice and flexibility as various competing technology providers serve the marketplace, and can interoperate while not assuming that multiple companies will mirror each other’s technology choices. Readers should, however, note the use of conditional statements (“has the potential to,” “can”) in describing the ability of FIM to provide the above benefits. The ability of FIM to successfully create such conditions is founded upon the implementation of appropriate policies related to the technology in use.

Among the factors that need to be considered, and which will be discussed further in this paper are: appropriate notice, choice and control options, data minimization, least-means access⁸, compliance, audit and oversight. While not yet a term of art, for the purposes of this paper, we will refer to such a responsible and accountable model as a Privacy and Trust-Enabled Federation.

7 An online Federated Identity Management system that, as of late, has been gaining user share is OpenID. Rather than registering a new ID with each newly visited site (which requires a new username and password as well as the disclosure of some level of personal information), a user is allowed to create an identity with a single “provider” (such as Yahoo, AOL, Verisign, etc.). When the user logs into the newly visited site, called the “relying party,” the provider is queried for authentication of the user; at this point, the user is redirected to the website of the provider, gives his or her password (if the user is not already logged into the provider), and is asked if he or she wishes to trust the relying party website. Upon agreement, the user is logged into the relying party’s site using his or her OpenID, thus creating for the user a single sign-in experience.

By taking the storage of identity information out of the hands of individual websites, OpenID allows users a higher degree of informational self-determination. This architecture allows for the possibility that rather than being forced to disclose personal information (and understand the associated usage policies) at each site that requires a log-in for full interaction, a single, trusted (by the user) identity provider can be chosen for all authentications.

8 Least-means access is a concept that bridges or combines two essential privacy concepts, namely least-privilege access and data minimization. Least-means access, in plainer language, is a principle that the least amount of information should be provided to meet a request and that should be further evaluated based on the access privilege of the requestor. This concept helps ensure that those requesting information have appropriate privileges to request, and an actual need for, the information they are requesting to accomplish a legitimate business purpose.

Along with the user, an FIM architecture typically contains (a minimum of) the following roles:

- **Service Provider (SP), or Relying Party (RP):** A Web application that provides a service to the user, but which has outsourced user authentication.⁹ This service thus “relies” on a third party to provide identity information. There will be multiple Service Providers within the FIM “ecosystem.”
- **Identity Provider (IP):** A website or service with which the user has established his/her identity.¹⁰ The IP provides identity verification services to the Service Provider, and may also be a central store of user information, to be distributed on a least-means access basis. There may be one or more IPs in the FIM ecosystem.
- **Discovery Service:** A means of finding an Identity Provider that is acceptable to both the User and the Service Provider; this could be as simple as a drop-down menu on the SP’s website.

While all the roles outlined above are part of an FIM architecture, they do not necessarily represent different entities. In fact, one should consider that a large enterprise may span the entirety of the architecture, with different organizations within the enterprise playing the various roles. Building upon these roles, Figure 1 describes a typical series of interactions within the FIM architecture:

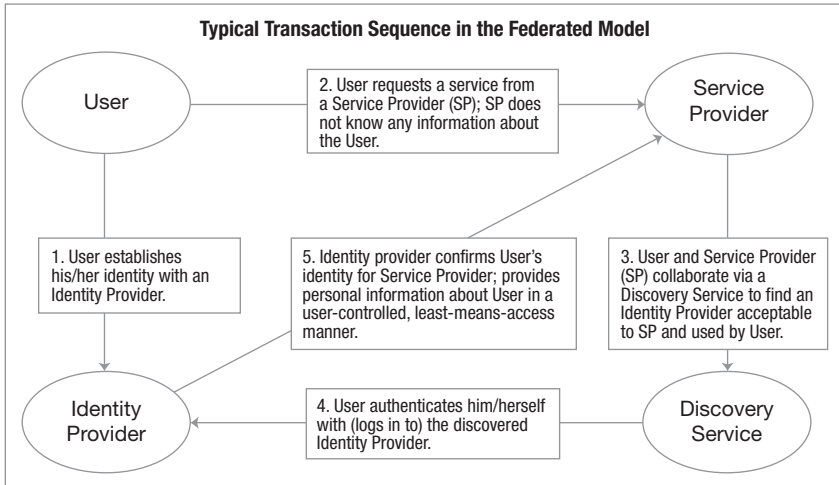


Figure 1: Federated Identity Management transaction sequence

9 Service Provider or Relying Party examples can include: public-facing Internet websites (such as an online store, blog, or community forum) and private enterprise systems (such as networked databases, procurement and human resource systems).

10 Identity Provider examples can include: public-facing Internet sites (such as banks, Internet Service Providers, or an online store) and private enterprise systems (such as corporate or partner directories).

By way of illustration, we will consider a Federated Identity Management system with which most readers will be accustomed: the banking industry's network of automated teller machines, or ATMs. Suppose a user has an account with Bank A (step 1 in Figure 1), but wishes to use an ATM belonging to Bank B, with which the user has no prior relationship (step 2). Once the user selects his/her desired transaction, the ATM must establish the user's identity; this is done by querying the discovery service integrated into the ATM network, which can recognize (via the card he/she has presented) that the user wishes to use Bank A for identity verification (step 3). Using the entered PIN number as verification of identity, the user effectively "logs in" to Bank A (step 4), which will confirm him/her as the account holder (step 5).

Suppose that the user's desired transaction is a withdrawal. Along with the log-in query, the ATM sends details of the proposed transaction to Bank A. In response, Bank A does not need to reply with detailed account information, such as name or address, the amount of funds currently being held (though this is often sent for inclusion on the user's receipt), whether the account has associated credit cards or loans, etc.; rather, the transaction can be completed with only a yes/no response to the ATM's inquiry (with some additional information sent for record-keeping purposes). The user is thus able to minimally disclose PII to the ATM, while completing secure (and potentially quite significant) financial transactions; this is the level of informational privacy possible within the FIM model.

2 Privacy

2.1 Essential Concepts

Information privacy relates to an individual's ability to exercise control over the collection, use, disclosure, and retention of his or her personal information. This notion is necessarily predicated on the provision of clear notice related to what information will be collected and how it will be used and/or shared. Personal information is any information, identifying or otherwise, relating to an identifiable individual. Specific PII may include one's name, address, telephone number, date of birth, age, marital or family status, financial status, e-mail address, etc. For example, credit cards, debit cards, social insurance/security numbers, driver's licences, and health cards contain a great deal of sensitive personal information. Moreover, it is also important to point out that almost any information, once linked to an identifiable individual, becomes personal information, be it biographical, biological, genealogical, historical, transactional, locational, relational, computational, vocational, or reputational. Hence, the definition of "privacy" can be quite broad in scope, rendering the challenges to privacy and data protection equally broad.

2.2 Why Are Privacy and Trust Essential in FIM?

Now that we understand what is meant by information privacy, what can be said of its importance to FIM? Why, in particular, should federations strive to achieve the status of Privacy and Trust-Enabled? The answer is as simple as “necessity.” End users, for instance, may already have an established comfort level with the policies and information uses of a particular company within a federation; this comfort does not necessarily extend, however, to other members of a federation. In order to encourage utilization of other federation services, these users need to create the same ease of mind with organizations across the federation, many of which will be wholly unfamiliar to them. Similarly, in order to encourage the internal strength and growth of the federation, enterprises that participate in this information ecosystem must feel that the established policies, procedures and technological rules are respected, regardless of how large the federated ecosystem. Thus, FIM is dependent on customers and federation members being confident in the ecosystem’s ability to facilitate the migration of trust between known entities and a broader range of organizations which the customer and/or federation member has little or no experience with, but that should be similarly considered trustworthy in the context of a given Privacy and Trust-Enabled Federation.¹¹

Privacy and trustworthiness may be more difficult to establish within a federation of multiple enterprises than within a single enterprise. In a lone enterprise, there is typically a common policy framework, technology implementation and user base; many tools exist, such as Privacy Impact Assessments, with which a company can demonstrate and delineate its data protection efforts. Across multiple enterprises, however, there will likely be many different policies, deployed technologies and types of users, all of which need to be both interoperable and consistent in the protections provided for shared data. Strong privacy measures undertaken by a single enterprise become meaningless if its data-trading partners do not have compatible measures; the policies and technologies of *all* federation members must satisfy the requirements of the trusting party.¹²

In order for a user’s trust to be maintained, individuals need to know that information about them that enters into this distributed web of systems will be handled according to established rules/frameworks and in conformance with promises made at the time of collection. For an ecosystem to be effective, rights to information must be apportioned across parties to meet legitimate needs, and users must be provided with appropriate controls related to their information. The majority of users, however, are neither capable of nor interested in micromanaging the ecosystem.

11 Note that the term “customer” is used broadly as a term that can encompass consumers, citizens and organizations.

12 An assurance that becomes easier in a least-means environment.

They will rely on their relationship with the organization that is their gateway/portal to the ecosystem, or with which they are ultimately connecting. Responsible organizations want to provide such assurances, both for building trust with individuals and for demonstrating compliance with various privacy regulatory environments.

Federated identity management presents a new challenge to privacy, in that transfers of personal information occur between organizations as well as between the individual and an organization. However, it is not the first technology to take the transfer of PII out of an individual's hands. In 1999, the IPC, along with the Dutch Data Protection Authority, looked at the concept of intelligent software agents.¹³ Such agents require access to the personal profile of the individual whom they serve, in order to complete tasks on behalf of a user without any direct supervision. Much as is the case of FIM, considerable user utility may be derived from the automation of routine tasks in this way. The user should not, however, have to sacrifice privacy in order to gain these benefits; appropriate controls on the functioning of the agents (limiting of scope, use of trusted sources, deployment of appropriate privacy-enhancing technologies, etc.) can, in fact, turn an agent from a privacy threat to a privacy protector. The same can be true for FIM.

Beyond the conceptual framework, practical aids are required to help system architects and implementers, as well as the people who develop the business and information use models, to enable trust at the design and development stage. To that end, this paper will aim to provide guidance for designers looking to build trust into systems, and describe useful tools to evaluate the level to which trust has been enabled and facilitate oversight and compliance.

2.3 How Can Privacy Be Enhanced by FIM?

Using Federated Identity Management to create a community of trust can provide significant benefits for the end user, as well as for participating organizations and oversight authorities. For an end user, the leveraged trust model permits him or her to authenticate across various different organizations with only a limited amount of information transmitted among organizations. Enterprises using this model only share information in a “least means necessary” manner, based on what is absolutely necessary for a given transaction. End users can also be assured of greater consistency of security requirements and policy conditions among the participating enterprises. Enterprises and organizations that participate in a community of trust have a basis for offering new kinds of services through a broader range of organizations, thereby providing greater value to end users through their trusted relationship. When applied across enterprises to an ecosystem, this combination of policies, practices, and tools, supplemented by contracts (where needed), creates

13 See the IPC publication *Intelligent Software Agents: Turning a Privacy Threat into a Privacy Protector*, April 1999, at: <http://www.ipc.on.ca/images/Resources/up-isat.pdf>

an overall privacy framework for the federation. Such a framework can then be compared against existing guidelines such as Fair Information Practices (FIPs). This will be beneficial to the design of a federation in that FIPs have been codified in many jurisdictions (Canada's *Personal Information Protection and Electronic Documents Act (PIPEDA)*, the OECD Guidelines on the Protection of Privacy and Transborder Data Flows of Personal Information, etc).

In our work promoting our idea of *Privacy by Design*¹⁴, the IPC has also established a framework concept for a Global Privacy Standard (GPS) for technology development.¹⁵ The GPS is not a technical standard per se, but rather a distillation of fair information practices and privacy principles, which are common to many of today's legal frameworks related to privacy and data protection. While the principles of the GPS do not represent a globally accepted norm, they serve as an effective preliminary guide to the introduction of important legal concepts and a "culture of privacy" into a federation. Below is a discussion of these principles, along with the practical ramifications for a Privacy and Trust-Enabled Federation member.

2.3.1 Consent

In a majority of legal systems that have addressed privacy, the individual's free and specific consent is required for the collection, use, or disclosure of personal (or sensitive) information, except where otherwise permitted by law. The "quality" of this consent is context dependent; clearer and more specific consent will likely be required as the sensitivity of the data rises. Consent is also not permanent, and may be withdrawn or revoked at a later date.

The "circle of trust" created by a Privacy and Trust-Enabled Federation allows member enterprises the possibility of collection and disclosure via the notion of "implied consent." This means that if the user is aware that services are being provided by a federation (i.e., clear and direct notice is given to the user), and not necessarily a single enterprise, then information can be shared for the *purposes specified at the time of collection* between trusted federation members, without the requirement of explicit user consent at each disclosure. This principle has been applied, for instance, in Ontario's *Personal Health Information Protection Act (PHIPA)*, which notes that health information custodians are permitted to assume implied consent (within the circle of care) to "collect, use, or disclose the information for the purposes of providing health care or assisting in providing health care to the individual," unless it is known that consent for such has been explicitly withdrawn.

14 "Privacy by Design" is a term coined in the '90s by the Ontario Information and Privacy Commissioner, Dr. Ann Cavoukian, in an effort to enlist the support of technology to protect privacy, rather than encroach upon it. By embedding privacy into the design of various technologies, and actually building it into the architecture of the technology involved, privacy is far more likely to be protected, instead of being viewed as an afterthought.

15 See Appendix 1 for the complete Global Privacy Standard.

2.3.2 Specified Purposes

As a second principle, organizations should specify the purposes for which personal information is collected, used, retained, and disclosed, and provide notice to the data subject at or before the time of collection.¹⁶ The specified purposes should be clear, limited and relevant to the circumstances.

2.3.3 Collection Limitation

Personal information should not be collected for collection's sake; rather, collection should be fair, lawful, and limited to that which is necessary for, and consistent with, the purposes specified to the data subject. As such, organizations should uphold the principle of *data minimization*, meaning that collection of personal data should be kept to a minimum. In addition, the design of programs, information technologies, and systems should begin with non-identifiable interactions and transactions as the default settings, and wherever possible, identifiability, observability, and linkability of personal information should be minimized.¹⁷

2.3.4 Use, Retention, and Disclosure Limitation

The use, retention, and disclosure of personal information should also be limited to the purposes identified to the individual, except where otherwise required by law. In addition to such a limitation, organizations within a Privacy and Trust-Enabled Federation should use and disclose data in a least-means access manner; that is, only the minimal data necessary (e.g., an age range, instead of an exact birthdate) for a specific transaction should be made accessible – this applies to both transfers of data between enterprise members and within the data-holding organization. Organizations should also retain personal information only as long as necessary to fulfill the stated purposes, after which time it should be securely destroyed.

2.3.5 Accuracy/Access

Neither organizations nor individuals should be satisfied with the use of inaccurate data within an enterprise or a federated ecosystem. In addition to the negative impacts on business, incomplete, out-of-date, or inaccurate data can have significant bearing on individuals, particularly in the case of qualifications required to access services in a federated ecosystem.

16 In some cases, information is collected as part of the technology handshake between systems – so things like the IP address of a computer and other related information may be captured before a notice of purpose may be provided. Some principles can thus be modified with the term “or as soon thereafter as practicable” to account for these automated collections related to system functionality.

17 There are, of course, transaction requirements that require identification from both a legal and business perspective; tax reporting, delivery requirements etc. We recall, however, the driver's licence example and highlight that even where identifiable information needs to be collected, concepts of data minimization should be applied.

As a principle, then, organizations should ensure, with varying degrees of exactitude depending on the context, that personal information is as accurate, up-to-date and complete as required for the specified purpose.¹⁸

In addition to organization-initiated data integrity checks, the issue of data accuracy can be addressed by the adoption of another privacy principle: access. This principle states that individuals should be provided with access to data collected about them, along with information about how that data is being/has been used and/or disclosed. Further, individuals should be provided a challenge mechanism, to address any perceived inaccuracies or deficiencies in the data. Allowing user access thus not only addresses a commonly held privacy standard, but it also brings the individual – in theory, one of the ultimate authorities regarding his or her personal information – into the solution of ensuring data accuracy.¹⁹

2.3.6 Security and Data Integrity

Organizations must assume responsibility for the security of personal information throughout its life cycle consistent with the international standards that have been developed by recognized standards development organizations. Personal information should be protected by reasonable safeguards, appropriate to the sensitivity of the information (including physical, technical and administrative means).

2.3.7 Accountability/Openness/Compliance

Finally, the collection of personal information entails a duty of care for its protection. Obligations related to all relevant privacy-related policies and procedures should be documented and communicated as appropriate, and assigned to a specified individual within an organization. When transferring personal information to third parties, organizations should seek equivalent privacy protection through contractual or other means. Further, in order to ensure the accountability of federation members, the principles of openness and transparency should be adopted. That is to say, information about the policies and practices relating to the management of personal information should be made readily available to the individuals whose data is being held.

Organizations should also establish compliance and redress mechanisms, and communicate information about them to the public, including how to access the next level of appeal. Compliance with privacy policies and procedures should be monitored, evaluated, and verified on an ongoing basis.

18 There is a complexity here, in that both access and accuracy can be difficult, if not impossible, to maintain down the chain of data custody. Such custodians must, however, ensure data security.

19 A number of enterprises have also found that self-service registration and maintenance of such information by the user not only increases accuracy, but also decreases the costs of operating and maintaining both the systems and the support functions.

The combination of these factors leads to the concept of accountability, in which information is appropriately secured and protected across both the scope of federated enterprises and the life cycle of the information. This concept not only forms the basis of Canadian privacy laws, but is found in the OECD Guidelines, and is the defining principle of the APEC Privacy Framework.

2.4 Role of Fair Information Practices and Global Privacy Frameworks

Frameworks such as GPS and Fair Information Practices are not, of course, intended to supersede legislation – an issue that *must* be considered within a federation, given the likely trans-border nature of data flows. Instead, such documents should be considered as guidance which can inform federations and federation members about essential privacy concepts, regardless of jurisdiction.

3 Risk Assessment

To effectively examine an information ecosystem, one must be familiar with the life cycle of the data contained therein. Information is collected, shared/used, re-freshed, or deleted, for a variety of purposes. The cyclical nature of the information life cycle must be supported by appropriate policies, practices, procedures, tools, and contracts. While the categories and principles of the life cycle are fairly constant, how they work, for what purposes they are used, and how they are supported varies from company to company, and ecosystem to ecosystem. These are dependent on the needs, types of information, sensitivity of information, policy choices and a number of their variables. Evaluating such life cycle considerations presumes a familiarity not only with the needs of the entities involved but also with the risks that those entities may face.

Thus, a foundational element of the Global Privacy Framework is a proper threat risk analysis.²⁰ Risk must be properly identified, minimized to the extent possible, and appropriately managed where it can't be eliminated. Recall that risks come in a broad range. Many considerations of risk are based on concepts of traditional compromise of systems – i.e., breach by an outsider. There are also risks that are harder to identify and control, arising from insider threats. Other risks include the potential for operational failure: from massive system failures to issues like failing

20 See Government of Canada publication *Threat and Risk Assessment Working Guide*, at: <http://www.cse-cst.gc.ca/publications/gov-pubs/itsg/itsg04-e.html>. Alternate risk assessment models include the National Security Agency's *INFOSEC Assessment Methodology*, at: <http://www.fountainheadcollege.edu/ia/nsa/iam.htm>, and the National Security Agency's *INFOSEC Evaluation Methodology*, at: <http://www.fountainheadcollege.edu/ia/nsa/iem.htm>, the US General Accounting Office (GAO) *Information Security Risk Assessment*, at: <http://www.gao.gov/special.pubs/ai99139.pdf>, or the National Institute of Standards and Technology (NIST) *Risk Management Framework*, at: <http://csrc.nist.gov/groups/SMA/fisma/framework.html>

to anonymize information sufficiently to prevent identification. These risks could be a matter of degree, or an issue of implementation/deployment.

A large part of appropriate risk mitigation and management is proper training, preparation, and incident response. No single system or ecosystem is foolproof, so training, preparation, and response are essential concepts because incidents invariably occur. A proper contemplation of the information life cycle includes these concepts. An entity that has taken the right preparatory steps will debrief incidents to see where improvements can be made, thus becoming a learning organization. In more advanced learning organizations, tests of processes and procedures related to incident response can provide much of the valuable learning, prior to an actual incident arising.

4 Federated Privacy Impact Assessment

For some time now, organizations with privacy compliance programs to protect privacy and the confidentiality of personal information collected use a tool known as a Privacy Impact Assessment (PIA). The PIA is one of many tools used to assist organizations to ensure that the choices made in the design of a system or process meet the privacy needs of that system, typically by way of a directed set of questions, based on privacy requirements. In some circumstances, it may be conducted alongside a threat/risk assessment, which is one of the inputs into assessing the overall privacy landscape. A PIA can be applied for the purpose of a systems requirement definition to a single project, or for the purpose of demonstrating regulatory compliance to an enterprise. A PIA must be considered in the context of the needs of the organization, the uses of the information, and the relevant fair information practices that are either required by regulation or otherwise adopted by the organization.²¹

“A privacy risk assessment should become an integral part of the design stage of any initiative. Once the risk to privacy is identified, then the necessary protections can be built in to minimize or ideally eliminate the risks.”²²

All PIAs should have a modular nature since most policies, governance frameworks, and systems are neither the purview nor expertise of a single person. A PIA will have a coordinator or point person within an organization, often the Chief Privacy Officer (CPO). The CPO should assemble the organizational team required to review

21 Organizations may develop policies and practices that interpret legal requirements; at times these policies and practices may exceed legal requirements, including where there are global practices that are in place despite the lack of any local regulatory or legislative requirements.

22 *The Privacy Payoff: How Successful Businesses Build Customer Trust*, Ann Cavoukian, Ph.D., and Tyler J. Hamilton, January 2002, p. 290.

and answer the PIA questions. In a corporate setting, that team would include representatives from technical support/customer service, security, marketing, and relevant lines of business. Optimally, the various owners/operators of the systems and other framework elements will have been consulted in the development of the PIA, and the PIA process will yield benefits to them as well. The PIA has two roles: one is assisting in privacy compliance, but the other, which has greater meaning for participants not responsible for privacy, is to be a building block of the information governance and risk management program of the company. The identification of the PIA in this manner will assist those disciplines not specifically concerned with privacy to better understand both the value of the review, its relevance to their job function, and the role it plays in adding value to the organization.

When organizations come together in a federated ecosystem, a new conceptual application of the PIA is required: the Federated Privacy Impact Assessment (F-PIA). The F-PIA differs from a traditional PIA in a number of ways. Most importantly, the F-PIA is designed to operate either within an enterprise (such as one where a number of different systems may be federated together) or across enterprises that have different needs and uses of information. An F-PIA must be designed to accommodate various starting points, from situations where most systems have already gone through some form of a PIA (in which case a focused gap analysis will be an effective starting point for the F-PIA), to more “green fields” settings where most of the systems requiring review have not undergone a PIA (in which case an F-PIA at the level of the federation can act as a template for initiating organization-level PIAs). While major elements of the F-PIA will seem familiar, they are designed to be applied in a more flexible fashion.

An F-PIA is designed to consider data “in motion” (that is, being transferred among various organizations), rather than data “at rest” (which is more relevant to a traditional PIA). The question of the exchange of information within an identity federation is in some ways more simple, and in some ways more complex, than questions asked in traditional PIAs completed by an organization. Simplicity can be found in that the nature of an identity federation will typically mean that the personal information that is involved in the federation is well specified and that data flows are documented. Each member of the federation has a defined role, or roles, such as Identity Provider (IP), Service Provider (SP), and Discovery Service (DS), and so on.²³ The type of personal information each role is entitled to should be well defined as part of the technical specifications of the federation. This fact by itself bypasses some of the most onerous aspects of a PIA within an organization, especially one with substantial amounts of unstructured data or data in legacy applications.

23 It is important to note that an organization may play more than one role in a federated ecosystem. As time and experience with these ecosystems progress, this merging of roles is becoming more commonplace.


There is, however, complexity in an F-PIA that does not exist in an organizational PIA. As the subject of an F-PIA is “data in motion,” it will, in some cases, mean that data will be passing between regulatory jurisdictions, between different industries, or across borders (PIAs should also take this consideration into account in the case of multinational corporations or the off-shore outsourcing of data). Thus, questions of technical accountability and custody of personal information arise that are not present in the single organization context. In many regulatory regimes, you may be able to outsource services, but you cannot outsource accountability. This means that every organization involved in a federation is a stakeholder, even if a particular data flow is not one in which it participates.

The F-PIA is also meant to be reusable. Due to the non-static nature of federations, it would be impractical to perform an ecosystem-wide privacy assessment for each added member or service. Such an undertaking is not expected. Instead, a new addition to a federation should be able to identify the role (or roles) that it will play, along with its corresponding responsibilities and requirements. In each of the models that follow, the opportunity exists to create, through an F-PIA, a set of privacy- and trust-enabling standards against which additions to the federation can be compared.

4.1 Nature of Federation

When an identity federation forms, establishes, or applies a set of privacy standards to be observed for the purposes of the federation, and elects to conduct an F-PIA, it will be faced with the question of how to coordinate the privacy policies, procedures, and practices of its disparate members. Federation can take a number of forms, ranging from a fluid and changing set of equal partners to a centralized group of subsidiaries, under the direction of a central body. The structure of the federation will establish where the authority to determine the privacy standards for that federation lies, subject to overriding legal constraints. The nature of federation will also aid those who undertake an F-PIA to determine the likely source of threats to such standards.


4.1.1 Collaborative Model



In the collaborative model, a group of founding members or member forms an entity that establishes the rules for the operation and governance of the ecosystem, as well as overseeing day-to-day control of the system. Described as the most complex of the models of federation, but with the greatest flexibility, this model is paradoxically likely to require the most rigid privacy rules and stringent F-PIA. These


controls are put in place to ensure that the indefinite membership and flexibility may not be exploited to extract PII for inappropriate uses. Assurances of minimum disclosure and strict technical enforcement of privacy guidelines will require audits and accurate user reporting to engender appropriate trust in verifiable privacy. The Governing Entity in the model will be the central authority for privacy compliance.

4.1.2 Consortium Model



In the second model, a small number of founders forms a consortium via a multi-party contract that sets the rules and governance for the ecosystem. Based on reasonably autonomous founders, the risk to privacy in the consortium model is that one or more of the founders may have a significantly different privacy model. With respect to the exchange of PII, the contractual agreement by which the federation is formed must be specific as to the common privacy elements. The F-PIA created for such a federation will need to be clear on the limits of the assertions that can be made for the consortium. It is very likely that the privacy assertions of the whole federation will be the “lowest common denominator” of the founders. Where consortiums develop from a common industry with a common expectation of practice, this may not present a significant bar, but in cross-industry consortia, this could generate friction.


4.1.3 Centralized Model



In the centralized model, a single founder sets the rules and governance for the ecosystem, and contracts individually with each other member. This approach provides the founder with a significant amount of control, and significantly less control to the other members. The centralized model ensures that data flows through,

or with the awareness of, the single founder, which implies that privacy assertions can be made and verified by that organization. This architecture also allows for the possibility of the single founder incorporating the data protections identified and afforded by the F-PIA to be contractually incorporated into the federation, in a highly uniform manner.

4.1.4 Service-Oriented Architecture



While complexity is increased in a service-oriented architecture (SOA), the analysis and impact assessment may actually be simplified. A useful way to think about an SOA is as a cloud of service elements that can be associated in a number of ways, either dynamically or in a directed fashion, to provide a service. The complexity is created by the potential for differing service elements to be joined together to create various desired solutions. In order to prove these results at a later time, there may need to be a record of the elements used in a given solution, as well as the characteristics of the elements since they may change periodically. In an F-PIA, an analysis related to an SOA architecture may be simplified by using the same modular approach. Taking cues from the design principles of SOA, F-PIAs can be assembled in a modular and reusable fashion as well by evaluating the various distinct modules and reusing the information gathered in an overall privacy evaluation of the inputs and ends of the service chain – the information that is accessed by the service, and the output of the service element used. The SOA environment should be evaluated on four parameters: **security** of the elements, **auditability** of elements, **access control** and **system oversight/accountability**. Once the elements have been evaluated, they can be reused in combination with greater confidence and lower overhead.

Regardless of the architectural model or legal form, in most cases, it is likely that there will be a mixed level of privacy practice across the federation. Some members, in particular, may have access to less, or less sensitive, personal data, and thus may need less elaborate protections and compliance procedures to provide appropriate safeguards. While there may be some temptation, for ease of implementation or for simplicity's sake, to "race to the bottom" and for the federation to adopt the "lowest" privacy standard of any of its members, this is likely to be counterproductive in the long run. This is particularly true if privacy is to be a source of competitive advantage

or the basis of trust between federating entities and their end users (citizens, consumers, and employees). Instead, member organizations must always ensure that they deploy data protections commensurate with the risks they face.

5 F-PIA Goals and Value Propositions

5.1 Data Subjects and Regulators

Besides the organization(s) involved in the federation, a number of other stakeholders exist who do not directly participate in the F-PIA, but who may utilize the results or rely on its process to safeguard their interests. Among the most important of these stakeholders are the data subjects, whether they be citizens, consumers, employees, or other individuals whose information is collected, processed, stored and shared by systems. Data subjects are beneficiaries of an F-PIA process to the extent that it appropriately ensures the security of the information and the implementation of related policies and procedures. Data subjects will be equally concerned that any promises made by entities within a federation are honoured and that appropriate complaint and recourse mechanisms exist.

Another major group of outside stakeholders is comprised of oversight and regulatory authorities. These stakeholders may include traditional privacy and information commissioners, relevant agencies responsible for oversight related to the protection of personal information, and private sector bodies that may be involved in trust oversight or evaluation. The F-PIA must consider a flexible way to address the needs of these stakeholders. PIAs, by their nature, may include secret and proprietary information that may not be appropriate to share in public documents. As such, the F-PIA process must consider ways to provide meaningful information related to the review without compromising either secret or proprietary information, or the kind of information related to configuration, which could be used to compromise systems security.

In this sense, the F-PIA must be considered in two parts. The first part is that which is relevant to the organizations involved and their compliance functions; the second part is a summary or redacted version of the first, which may need to be provided to oversight organizations or, where appropriate, made available to the public. The reduced information made accessible as an outcome of an F-PIA should not be viewed as a method of hiding deficiencies, but rather an opportunity to ensure that information that could compromise security, or should otherwise be protected by the organization, is not released. Under appropriate conditions, of course, oversight authorities may be able to request or require the production of more detailed information to meet their legal obligations.

5.2 What are the Goals of an F-PIA?

There are four primary goals to be achieved through an F-PIA.

Goal 1: To provide an opportunity for members to discuss, develop, and codify a federation's privacy policies.

First, it is recognized that privacy policies will vary by federation. These policies should address fair information practices as appropriate to the contextual application of the federated ecosystem and the regulatory requirements to which it may be subject. In addition, the policies should recognize that the person whose data is being processed should be provided with appropriate choice and control over both who has access to, and what can and cannot be done with, their data (with allowances being made for overriding factors such as court orders or medical emergencies). However, regardless of particular policy choices, individuals must be convinced of the veracity of a federation's claims of data protection in order to create a trusting relationship. Thus, the second goal of an F-PIA:

Goal 2: To demonstrate that privacy policies, as defined by the members of the Federation, will be met.

Most privacy policies require some degree of data minimization and enable legal requirements of choice or consent. Thus, an F-PIA of a system that undertakes consent-based collection of personally identifiable information must ensure that consent is properly addressed, as either open-ended or specific, opt-in or opt-out, depending on the requirements. An F-PIA evaluation of a data-minimizing policy would likewise require that only the minimum possible personally identifiable information is collected. An F-PIA would use similar analysis frameworks to address all covered elements of the policies and legal requirements, such as determining the extent to which data is shared, the uses to which it is put, and the length of time that it is retained.

Effective privacy protection, however, requires a number of system design elements to be in place, irrespective of the particular privacy standards that are applied to a system. Without a doubt, up-to-date, robust security mechanisms must be in place to ensure that access to data can be reliably restricted to only those who have an established right to the data, as established by privacy policies. Thus, the third goal of an F-PIA:

Goal 3: To demonstrate that an appropriate technological architecture is in place to prevent, to the extent possible, accidental or malicious violations of privacy policies.

One must recall that security does not equal privacy; security is critical to privacy, but it is a contributory factor. To perform this function, single demonstrations of technological privacy protections and security measures are not sufficient. An F-PIA should be an iterative and ongoing process. Privacy is not a momentary commitment, nor should an F-PIA be a "box ticking exercise." It is an ongoing

obligation to actively meet the needs of organizations, promises made to individuals, and regulatory requirements. It behooves the federation to create a goal of the highest achievable privacy standards, and to revisit this goal on a regular basis. An F-PIA should be conducted for any new system or program, as well as at the point of any substantial change to systems or programs handling PII. Since an identity federation exchanges identity-related information, any substantial change in the federation or its data flows should also be accompanied by a revisiting of the F-PIA, as should any major privacy breach. The F-PIA is a living document, a tool always to have at hand either for use or revision.

Ensuring that F-PIAs are completed with appropriate candour and resources brings us to the fourth goal of an F-PIA:

Goal 4: An F-PIA should benefit all parties who complete, use, and rely on an F-PIA.

In this paper, we explored various stakeholder interests in an F-PIA. Many people who design PIAs focus only on the ultimate regulatory object of the PIA process without considering the value it can and should return to those completing the PIA. This is an especially important concept in an F-PIA, where a top-down directive within an organization is not the main motivating factor. The question for the F-PIA architect is, how to make sure that the questions and format provide useful information about a system to the designers and users of the system beyond its ability to comply with stated requirements. Often the benefits are the clarification of objectives and obligations, as well as the interaction between technology and policy across the system. Thus, architects should be aware that some aspects of the results of the F-PIA may need to be shared across entities to enable them to better understand the systems design and interrelation.

Ultimately, though, it is individual consumers who benefit most from this process, through the privacy protections afforded them. This benefit then trickles down to all federation members through competitive advantage, increased consumer confidence, and broadened consumer usage of federation services.

5.3 F-PIA Framework

Once the privacy requirements have been established, and prior PIAs have been integrated, the F-PIA becomes a matter of conducting an assessment across the entities in a federation. This is another iterative process, where strategic steps are iterated at finer levels of granularity as the F-PIA winds its way from high-level objectives to concrete determinations at the level of data and individual procedures. The key structural elements of a functioning F-PIA are the data itself, policies and procedures, technology and systems, and accountability. By creating an assessment based on appropriate fair information practices (e.g., Global Privacy Standard) that accounts and reports on these elements, an identity federation will be capable of providing assurance to all stakeholders regarding privacy-related issues.

6 Questions You Should Be Asking

The elements that need be examined in an F-PIA can roughly be divided into the following three categories: the Information Life Cycle, Organizational Principles, and Implementation. In this final section, sample areas of inquiry will be given for each. Please note, though, that the following questions are not meant to be comprehensive, nor necessary for inclusion in the F-PIA. Rather, they are presented as a means of suggesting the types of questions that should be asked in regards to privacy and security standards within a federation.

6.1 Information Life Cycle

We previously stated that it is important to create a “culture of privacy” within an organization, in order to create a trusting relationship with users. The questions asked regarding the Information Life Cycle attempt to examine this culture. In particular, a federation should consider its treatment of personal information, whether it is collected for necessary purposes, and whether its dissemination is ultimately decided upon by the individual involved – that is, the federation should contrast its practices with those prescribed by generally accepted privacy standards. Areas that should be explored within this topic may include:

- 1 **Appropriate Notice** – Is the individual whose personal information is being transferred aware of the transfer?
- 2 **Appropriate Specification** – Are the federated parties appropriately aware of the limitations related to the collection, use, sharing, and retention of information?
- 3 **Appropriate Consent** – Can transfers of personal information be *appropriately* linked to a user’s consent or choice?
- 4 **Appropriate Control** – Does the user have appropriate control over the transfer of his or her personal information?
- 5 **Data Minimization** – Do federation members collect the minimum amount of personal information necessary?
- 6 **Least-Means Access** – Do federation members transfer/access only the personal information needed to complete a particular transaction?
- 7 **Compliance, Audit and Oversight** – Is there an oversight body, or auditing or compliance mechanism, to ensure that privacy policies are met?
- 8 **Reporting** – Is there sufficient documentation of policies and procedures to help demonstrate compliance?

6.2 Operational Principles

An examination of the operational principles of the federation should fully describe the philosophy of interactions both among federation members, and when communicating with persons collecting data from them. Here, a focus on clarity is required: it must be shown that each member understands what is required of it within the federation. It is at this level, in particular, that the true complexities of federation may arise.

Questions provided for guidance in examining operational procedures include:

- 1 **Structure/Role Assignment** – Are the roles of all federation members clearly understood and transparently defined? Do federation members know their responsibilities and obligations?
- 2 **User Understanding** – Are the names or types of members of the federation and their roles made clear to the user?
- 3 **Identity Management at the Ecosystem Level** – Do Service Providers have the capacity to link a user’s profile across services, in the absence of user authorization? This may be the case if a Service Provider serves a dual role as the Identity Provider. [This topic goes to the ability of federated identity formats to enable appropriate sharing limitations.]
- 4 **User Involvement** – How does the federation protect against account linking, traffic and analysis? How does the federation encourage user involvement in defining controls?
- 5 **Worst Case Scenario** – Has a “disaster” scenario been considered, including steps to be taken to notify users and minimize any damage that may have resulted?

6.3 Implementation

An F-PIA should consider the various elements of the technical implementation of the Federation. Beginning with the design and architecture of the system, an F-PIA should include undertaking an assessment of the flows of information and how the technology is configured to ensure the privacy goals of the community of trust, or what we are calling federation. Drawing from the OECD Security principles and Liberty Alliance’s best security practices,²⁴ the following framework might be followed when developing detailed questions in this area:

24 Liberty Alliance paper, “Privacy and Security Best Practices,” Version 2.0, November 12, 2003, at http://www.projectliberty.org/liberty/strategic_initiatives/privacy_trust_security

- 1 **Awareness** – Are federation members aware of the need for information and network security, and the steps they can take to enhance security?
- 2 **Accountability** – Are federation members accountable for information security, to the extent appropriate to their role?
- 3 **Response** – Is there a response action plan in place, so that federation members can co-operatively prevent, detect, and respond to security incidents?
- 4 **Ethics** – Do participants understand that their own action or inaction may harm other federation members?²⁵
- 5 **Risk Assessment** – Have all federation members individually, and at the level of federation, completed risk assessment and minimization processes?
- 6 **Security Design and Implementation** – Is security designed as an essential element of the information systems?
- 7 **Security Management** – Does the federation have a comprehensive approach to security management?
- 8 **Reassessment/Learning** – Do the federation and federation members have a schedule for reassessing security measures and making modifications as appropriate, including reassessment after incidents or operational failures?

In addition to inter-federation security measures, technical questions regarding common security threats at the user-federation member transaction level must be addressed. These threats may involve denial of service, message replay, spoofing, brute force, or many other common forms of online attack. Sample questions that a federation, and each of its individual members, may wish to ask include:

- 1 Are user interactions (beyond the log-in process itself) authenticated? If not, what alternative measure is used to prevent session hijacking?
- 2 Will session tokens be used? If so, what measures are in place to prevent message replay?
- 3 Have authentication measures been evaluated to ensure that they are appropriate to the nature and sensitivity of the information?

Again, these questions are not meant to be comprehensive, but instead are meant to provide examples of issues that must be addressed when analyzing privacy and security measures within a federation.

25 Often couched as a democracy principle when applied by government organizations, all F-PIAs should have an objective ensuring that the security of information and networks is compatible with the essential values of a free society (such as free exchange of ideas, openness, transparency).

7 Next Steps

One of the concepts that the Information and Privacy Commissioner of Ontario originated and has been a leading voice for is *Privacy by Design* – advancing the concept of building privacy directly into technology, as part of the design and deployment process. This office strongly believes that there are great potential benefits for both consumers and organizations in the deployment of Federated Identity Management; however, these benefits can only be fully realized in the context of a Privacy and Trust-Enabled Federation. We also believe that one of the most important and effective tools for demonstrating the adoption of *Privacy by Design* is the Federated Privacy Impact Assessment (F-PIA).

Having examined the above material, what would be the next step for a federation? It would be the development of a formal F-PIA. This paper is only intended to serve as a guide – organizations and federations must use it, along with the numerous PIA development tools currently in existence, to create measurable standards against which privacy and trust measures can be compared. Ultimately, one must remember that in this process, a zero-sum game is not at play – functionality does not need to be traded off for privacy. Rather, building in privacy and trust in a positive-sum manner creates a win-win scenario, in which both consumer and supplier are the beneficiaries of a robust information ecosystem. Unnecessary trade-offs should become a thing of the past.

Appendix 1: Global Privacy Standard

- 1 **Consent:** The individual's free and specific consent is required for the collection, use or disclosure of personal information, except where otherwise permitted by law. The greater the sensitivity of the data, the clearer and more specific should the quality of the consent required. Consent may be withdrawn at a later date.
- 2 **Accountability:** Collection of personal information entails a duty of care for its protection. Responsibility for all privacy-related policies and procedures shall be documented and communicated as appropriate, and assigned to a specified individual within the organization. When transferring personal information to third parties, organizations shall seek equivalent privacy protection through contractual or other means.
- 3 **Purposes:** An organization shall specify the purposes for which personal information is collected, used, retained, and disclosed, and communicate these purposes to the individual at or before the time the information is collected. Specified purposes should be clear, limited and relevant to the circumstances.
- 4 **Collection Limitation:** The collection of personal information must be fair, lawful and limited to that which is necessary for the specified purposes.

***Data Minimization** – The collection of personal information should be kept to a strict minimum. The design of programs, information technologies, and systems should begin with non-identifiable interactions and transactions as the default. Wherever possible, identifiability, observability, and linkability of personal information should be minimized.*

- 5 **Use, Retention, and Disclosure Limitation:** Organizations shall limit the use, retention, and disclosure of personal information to the relevant purposes identified to the individual, except where otherwise required by law. Personal information shall be retained only as long as necessary to fulfill the stated purposes, and then securely destroyed.
- 6 **Accuracy:** Organizations shall ensure that personal information is as accurate, complete, and up-to-date as is necessary to fulfill the specified purposes.
- 7 **Security:** Organizations must assume responsibility for the security of personal information throughout its life cycle, consistent with the international standards that have been developed by recognized standards development organizations. Personal information shall be protected by reasonable safeguards, appropriate to the sensitivity of the information (including physical, technical, and administrative means).

- 8 **Openness:** Openness and transparency are key to accountability. Information about the policies and practices relating to the management of personal information shall be made readily available to individuals.
- 9 **Access:** Individuals shall be provided access to their personal information and informed of its uses and disclosures. Individuals shall be able to challenge the accuracy and completeness of the information and have it amended, as appropriate.
- 10 **Compliance:** Organizations must establish complaint and redress mechanisms, and communicate information about them to the public, including how to access the next level of appeal. Organizations shall take the necessary steps to monitor, evaluate, and verify compliance with their privacy policies and procedures.

Please see *Creation of a Global Privacy Standard* at:

<http://www.ipc.on.ca/index.asp?navid=46&fid1=575>

References

Online Privacy

7 Laws of Identity: The Case for Privacy-Embedded Laws of Identity for the Digital Age (October 2006) at:

<http://www.ipc.on.ca/index.asp?navid=46&fid1=471>

Creation of a Global Privacy Standard (November 2006) at:

<http://www.ipc.on.ca/index.asp?navid=46&fid1=575>

Privacy in the Clouds: Privacy and Digital Identity – Implications for the Internet (May 2008) at: <http://www.ipc.on.ca/index.asp?navid=46&fid1=748>

Privacy and the Open Networked Enterprise (December 2006) at:

<http://www.ipc.on.ca/index.asp?navid=46&fid1=576>

Privacy and Security

A View from 2018: A Glimpse of the Internet Future (June 2008) at:

www.biac.org/members/iccp/mtg/2008-06-seoul-min/Final_View_from_2018_ICCP_Chair_Paper.pdf

Intelligent Software Agents: Turning a Privacy Threat into a Privacy Protector. Result of a joint project of the Office of the Information and Privacy Commissioner/Ontario and the Registratierkamer, The Netherlands. April 1999.

<http://www.ipc.on.ca/index.asp?navid=46&fid1=316>

Privacy and Security Best Practices (version 2.0, November 2003) at:

http://www.projectliberty.org/liberty/strategic_initiatives/privacy_trust_security

Cavoukian, Ann, Ph.D., and Hamilton, Tyler J., *The Privacy Payoff: How Successful Businesses Build Customer Trust*, pp. 290, January 2002.

Risk Assessment

Information Security Risk Assessment (August 1999) at:

<http://www.gao.gov/special.pubs/ai99139.pdf>

INFOSEC Assessment Methodology at:

<http://www.fountainheadcollege.edu/ia/nsa/iam.htm>

INFOSEC Evaluation Methodology at:

<http://www.fountainheadcollege.edu/ia/nsa/iem.htm>

Risk Management Framework (August 2008) at:

<http://csrc.nist.gov/groups/SMA/fisma/framework.html>

Threat and Risk Assessment Working Guide (November 2005) at:

<http://www.cse-cst.gc.ca/publications/gov-pubs/itsg/itsg04-e.html>