June 2, 2014

Trustmarks and Privacy

Annie Antón, Douglas M. Blough, E. Anwar Reddick, and *Peter Swire*[*]
Georgia Institute of Technology

This paper addresses the potential impacts on privacy, both positive and negative, of adopting trustmarks in electronic systems and networks where identity provision and verification are important components. Trustmarks are modular certifications that can be accepted across a wide variety of systems and communities. Trustmarks have been proposed as a key component within the "Identity Ecosystem", which is being fostered by the National Institute of Standards and Technology under the National Strategy for Trusted Identities in Cyberspace (NSTIC) Program [1]. As part of this program, the Georgia Tech Research Institute (GTRI) is developing a "trustmark framework", which is centered around three principles:

- Trustmark framework components should be machine readable;
- The trustmark framework should be flexible enough to support a wide range of Identity Ecosystem participants and requirements; and
- Enabling the verification of the authenticity and integrity of trustmarks is paramount.

For the purposes of discussion in this paper, we use the GTRI trustmark framework as an illustrative example and we therefore treat these three principles as a given.

The privacy analysis of the trustmark framework is greatly simplified because the framework focuses on federated identity management practices and other actions by identity providers and other institutions, rather than on the personally identifiable information (PII) of individuals. For instance, a group of trustmarks may indicate that a particular state or local government agency complies with a standard for issuing a credential, such as the Federal Identity, Credential, and Access Management (FICAM) program (http://www.idmanagement.gov/) or the National Identity Exchange Federation (NIEF, https://nief.gfipm.net/). For privacy purposes, the key fact is that trustmarks facilitate actions by the state or local agency, such as to show that employees of that agency are governed by a set of security, privacy, and other rules. Trustmarks primarily involve information about an organization's activities, such as its compliance with security requirements. The involvement of individuals in the trustmark framework is quite limited, and is discussed in detail later in the paper.

Section 1 of this paper provides background on the GTRI trustmark framework. Section 2 examines potential privacy risks that could arise with adoption of the trustmark framework. These risks are low because of the focus on actions by organizations rather than individuals. Section 3 examines potential privacy benefits that could arise with adoption of the framework. Although many of the initial trustmarks

[*] This work was performed under financial assistance of Award 70NANB13H189 from U.S. Department of Commerce, National Institute of Standards and Technology.

are likely to focus on security and interoperability, it is also possible to develop trustmarks that indicate compliance with privacy standards. To the extent these privacy trustmarks are developed and deployed, they can assist in scaling privacy protections on the Internet.

1 Introduction to Trustmarks and the Trustmark Framework

This section introduces the GTRI trustmark framework in preparation for discussions of related privacy issues in Sections 2 and 3. We do not provide a comprehensive description of the trustmark framework; see the GTRI trustmark website [1] for more detail.

1.1 Trustmark Framework Overview

The goal of the trustmark framework is to facilitate federated identity and attribute management, in other words the reuse of digital identities (herein "identities") and associated attributes, at Internet-scale. By reuse of identities at Internet-scale, we mean the commoditization of access to reliable identity information similar to the existing commoditization of Internet access. Identity reuse requires trust between entities that assert identities/attributes (identity/attribute providers) and entities that rely on such assertions (relying parties). The rules and requirements for establishing such trust comprise an identity trust framework (herein "trust framework"). The requirements of a trust framework may be explicitly or implicitly stated, and may encompass many dimensions such as identity assurance, privacy, security, technical interoperability, business-level identity requirements, legal rights, responsibilities, liabilities, and indemnification.

To date, reuse of identities for transactions that require non-trivial levels of identity assurance has been limited to occurring within small communities of interest (COIs) of identity/attribute providers and relying parties, where each community defines its own trust framework. This limitation exists because most trust frameworks are monolithic and expressed in non-standard formats, which makes the process of understanding and adhering to a trust framework costly and lengthy for enterprises.

The GTRI trustmark framework engenders the componentization and standardized expression of trust frameworks, which improves their comprehension and allows for reuse of their components. This quality of the trustmark framework facilitates cross-community reuse of identities that approaches Internet-scale, by reducing the burden of organizations' adherence to multiple trust frameworks. Modular, formalized trust framework components are called trustmarks (see Figure 1). The conglomerate of COIs that wish to reuse identities at Internet-scale is called the *Identity Ecosystem*.

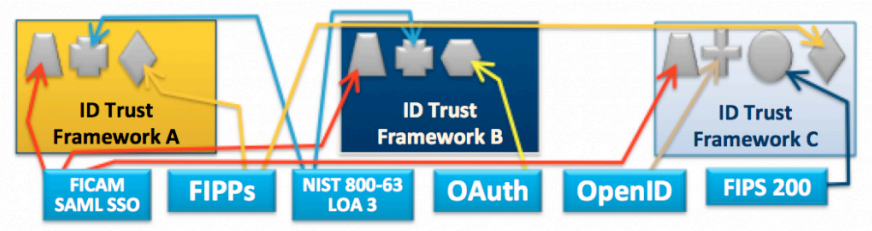

Figure 1: Trust Frameworks Based on Modular Trust Components, or "Trustmarks"

The notion of the GTRI trustmark framework has some similarity to a public key infrastructure (PKI) [10] (see Figure 2), which employs third party certificate authorities to enhance trust in an entity's claim to an identity. We now define the components of the trustmark framework (also see Figure 3).

Trustmark Framework Concept	Analogous Concept from PKI
Trustmark	Certificate
Trustmark Provider	Certificate Authority
Trustmark Recipient	Subscriber
Trustmark Relying Party	Certificate Relying Party / Audience
Trustmark Policy	Certificate Policy
Trustmark Agreement	Subscriber Agreement
Trustmark Defining Organization	ITU (Agency that defined X.509)
Trustmark Definition	IETF RFC 5280 (X.509 Spec)
Trust Interoperability Profile	List of Trusted Certificate Authorities

Figure 2: Parallels between the Trustmark Framework Concept and the PKI Concept

A trustmark is an artifact that is a statement of its possessor's conformance to a well-scoped set of trust and/or interoperability requirements, and is analogous to a PKI certificate. A trustmark provider (TP) issues a trustmark to a trustmark recipient (TR) based on a formal assessment process. A trustmark is issued under a trustmark policy, which defines the applicability and scope of the trustmark. A trustmark is issued subject to a trustmark agreement, which defines the legal rights and responsibilities of parties to the issuance of a trustmark. A trustmark definition (TD) defines the meaning and conformance criteria of a trustmark. A TD also includes the formal assessment process for the trustmark that TPs must follow. A TD is developed and maintained by a trustmark defining organization, which represents the interests of one or more COIs. Possession of a set of trustmarks by a TR is required by a trustmark relying party (TRP), which treats the trustmarks as 3rd-party-verified evidence that the TR meets the trust and/or interoperability requirements for identities that are set forth in the associated TDs. Note that a TRP is not necessarily a relying party. A TRP relies on trustmarks issued to a TR, while a relying party relies on assertions of identity and attributes provided by identity/attribute providers. A COI, an organization, or an individual may be a TRP and may publish and/or maintain a trust interoperability profile (TIP), which expresses the trust and interoperability criteria of the publisher as a set of trustmarks that TRs must possess. TRPs are not required to publish their TIPs, but they may do so to aid TRs in the automated discovery of TRPs that have requirements that are met by the TRs. Also, the rules and requirements for governing participation in a trust framework can be expressed as a TIP.

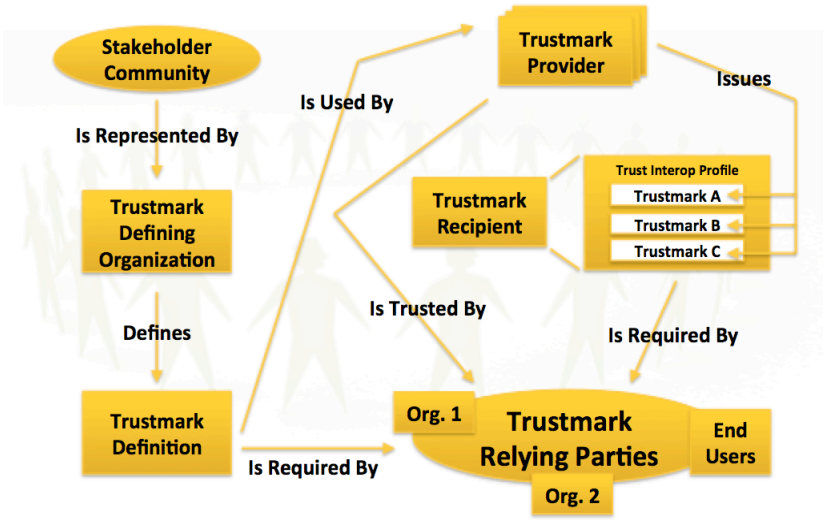

Figure 3: The Trustmark Framework Concept Map

Trustmarks, trustmark definitions, and trust interoperability profiles are self-contained, machine-readable artifacts that can be cryptographically signed to ensure their authenticity and integrity. Furthermore, these artifacts can be placed into online registries to support their discoverability. Trustmark framework artifacts can be published, discovered, and exchanged in order to configure trust relationships between various parties. These transactions are considered to be done at "trust-time". They provide a trust foundation for, and are distinct from, business transactions (such as conducting federated authentication) that are done at "run-time".

The trustmark framework can be adapted to meet the needs of a wide range of COIs. Examples of COIs that would benefit from reuse of trusted identities include varying levels of government (such as federal, state, tribal, and municipalities), varying functions of government (such as law enforcement or emergency management), industry sectors (such as health care and banking), and communities that cater to private individuals.

1.2 Trustmark Framework Example Use Case

As an example of the use of the trustmark framework, consider Agency-A and Agency-B, which are two law enforcement agencies that are geographically distant and have never before shared trust metadata or business information with each other. Suppose an officer of Agency-A wishes to access an information system operated by Agency-B for the purpose of obtaining information about an individual that has been suspected of committing a crime in the jurisdiction of Agency-A and is known to be currently in the jurisdiction of Agency-B. The officer possesses an electronic authentication credential managed by Agency-A. The goal is to allow the officer to authenticate to Agency-A's identity provider and have it issue an assertion, which attests to the authentication event, identity, and attributes about the officer (in this case, the identity provider is also serving as an attribute provider), to Agency-B's information system (in this case, the information system is a relying party), which will then authorize the officer's access to the data in question. For this goal to be achieved, Agency-A will need to have a trust relationship with Agency-B.

Agency-B will want assurances that Agency-A manages the identities and authentication of its officers in a trustworthy manner, and may adopt the FICAM requirements for identity level of assurance (LOA) 3 [7]. A trustmark defining organization acting on behalf of the FICAM program can define appropriate trustmark definitions for LOA 3 (GTRI is developing a preliminary set of such trustmark definitions, see [1]). Agency-B can publish a trust interoperability profile that references these trustmark definitions. This profile will require Agency-A to obtain the associated trustmarks from trustmark providers trusted by Agency-B. In this case, Agency-A is acting as a trustmark recipient and Agency-B is acting as a trustmark relying party.

In addition, Agency-A will want assurances that Agency-B maintains audit records of run-time business transactions in a secure manner. Agency-A can act as a trustmark relying party and publish these requirements in its own trust interoperability profile. Agency-B will have to obtain the appropriate trustmarks as a trustmark recipient, from trustmark providers trusted by Agency-A. In reality, Agency-A and Agency-B will likely have additional trust requirements that were not discussed in this example.

Once all of the necessary trustmarks have been collected and verified by both agencies, the agencies have a mutual trust relationship established. After this "trust-time" interaction, requests by an officer such as described above, can take place seamlessly at run-time. We emphasize that the trustmark framework is involved in this process only at trust establishment time, but not at run time.

2 Privacy Risks in the Trustmark Framework

This section examines privacy risks to individuals in the trustmark framework. Much of the focus of the NSTIC program is on authentication and identification of individuals, and there can be significant privacy risks associated with these activities. By contrast, the trustmark framework is not an authentication or identification approach. This section briefly summarizes privacy risks in authentication and identification, and then contrasts those risks with the operation of the trustmark framework. Notably, trustmark frameworks generally involve information about organizations or individuals acting on behalf of organizations, but in almost all cases do *not* involve the transfer of personally identifiable information (PII) about customers or users. There are possible privacy risks where individuals act as self-agents, i.e., on behalf of themselves. These situations do not occur often in the context of trustmark frameworks but they can happen. The primary role an individual self-agent can play in the framework is as a trustmark relying party; this scenario is discussed in more detail in Section 2.3.

Theoretically, an individual may also be a trustmark recipient. However, to our knowledge, no such use cases for trustmark frameworks have as yet been identified. This is because the task of attesting to user attributes is largely assigned to third-party identity providers and attribute providers in the identity ecosystem. However, if such use cases do arise, then it would be important to perform a detailed privacy analysis on them.

2.1 Privacy Risks in Authentication and Identification

Privacy protection is emphasized in the National Strategy for Trusted Identity in Cyberspace because many methods of authentication and identification have significant privacy risks. The existence of these privacy risks is a reason for detailed discussion of privacy considerations by each grantee of the NSTIC. In sharp contrast, the trustmark framework has very low privacy risks. Understanding the privacy risks of

authentication and "trusted identity" more generally enables a clearer understanding of why the trustmark framework has very low such risks.

One risk is that supposedly secret information, such as a Social Security number (SSN), becomes known over time to more parties. The use of SSNs for establishing identity, particularly for the task of creating new user accounts, illustrates the problem. SSNs have some good properties for identification – precisely one SSN is issued per individual, and the government has long-run reasons based on running the tax system to ensure that the one-to-one match between individual and SSN persists over time. In earlier years, the SSN had an additional advantage – it was known to the individual but not to most other parties. Over time, however, the usefulness of the SSN (alone) as an identification method has decreased. The SSNs of many people are posted on the Internet, such as on real estate records, and there are so many databases that contain an individual's SSN that we cannot have confidence that a person who states an SSN to a relying party is actually the correct person. In addition, each time an individual uses the SSN with a new party, such as a new health provider or merchant, yet one more (potentially insecure) database contains the SSN.

This privacy risk due to the revelation of secrets is also common to many authentication approaches, although not to the trustmark framework. The classic ways to authenticate are what you know (such as a password), who you are (such as a biometric), and what you have (such as a driver's license). Biometric approaches to authentication, such as fingerprints, can contain substantial privacy risks [2]. One risk occurs because the individual providing the fingerprint reveals the fingerprint to another party, who, when checking the fingerprint, essentially has the opportunity to take an accurate picture of the fingerprint. Security expert Bruce Schneier has explained that, once another party has an accurate picture, a laser printer and low-cost gelatin enable a fake fingerprint for less than $10 [2].

Similar problems can exist with something the individual has, such as a driver's license. In the era of inexpensive digital photographs, each time that an individual hands the driver's license to a relying party, there is the possibility that the relying party takes a high-quality photo. At a minimum, the relying party can learn the information on the license, such as license number or date of birth. That information, in turn, can assist the relying party to engage in identify fraud. More broadly, the high-quality photo can provide the basis for a fake driver's license. Governments have sought to "harden" driver's licenses and make them more difficult to fake, but the act of handing the license to a relying party means that a potentially insecure party gets the information on the license.

In sum, in many scenarios, especially those that require a high level of identity assurance, individuals provide identifying information about themselves to relying parties and registration authorities [10]. Each of these entities becomes a potential source of a data breach or identity fraud. Although the trustmark framework does not have this sort of privacy risk, the longstanding privacy risks in many systems thus provide compelling reasons to assess privacy risks in connection with initiatives developed under the National Strategy for Trusted Identities in Cyberspace.

2.2 General Privacy Risks in the Trustmark Framework

The authentication methods just discussed each required revealing of PII, such as the SSN, the fingerprint, or the information on the driver's license. U.S. privacy laws apply to PII or similar concepts but not to other information, such as information about organizations. For instance, the Privacy Act of

1974 applies to PII, while the HIPAA medical privacy rule applies to "protected health information" and the Gramm-Leach-Bliley Act applies to "nonpublic personal information." There can be significant privacy risks in federated identity management systems. However, this paper examines privacy risks in the GTRI trustmark framework. In contrast to many other systems that facilitate federated authentication and identity management, there are very low privacy risks in deploying the trustmark framework, primarily due to the lack of transfer of PII in the use of trustmarks by organizations to establish trust.

The trustmark framework facilitates wide-scale trust at the institutional level among entities such as identity providers, attribute providers, and relying parties. Each trustmark provides a modular credential, stating that an entity meets a prescribed set of identity trust and/or interoperability requirements. For the example of multiple police departments, the police department seeking information for an investigation (the trustmark recipient) contacts the other police department that has the relevant information (the trustmark relying party). The trustmark framework enables the two organizations to know that the other meets the security and privacy standards that each requires in order to conduct business. Thus, the trustmark framework artifacts, in this use case, contain information primarily about the organizations and not about individual users within those organizations. This example shows why the trustmark framework has categorically different and far less severe privacy risks than actual runtime information exchange mechanisms such as authentication mechanisms.

We have identified several instances where information about individuals can appear in the trustmark framework or where there are other privacy implications. First, contact information might be included in trustmarks. For example, the trustmark relying party might need contact information for the trustmark provider if they have questions about details of the trustmark. Similarly, the trustmark relying party might want contact information for the trustmark recipient if they have questions related to the trustmark recipient's security and privacy policies and practices. This could reveal the identity of individuals within an organization such as the Certification Officer for a trustmark provider or the Chief Security Officer for a trustmark recipient. This identification of an employee, however, is entirely different than transfer of PII about an individual whose data is contained within the IT system. Identification of each employee is a common element of audit and accountability measures in IT systems. For instance, the HIPAA privacy and security rules require a single sign-on for each employee of a hospital or other covered entity; in that way, if there is a violation of the rule, it is far easier to hold the violator accountable. Under U.S. law, the IT system is considered to be property of the organization, and individual employees (with limited exceptions) have no legal right to privacy for their actions on their employers' computers. So, from a legal standpoint, revealing employees' identities is not a privacy risk.

Second, widespread adoption of the trustmark framework will facilitate more interactions between different parties, because of the capability to identify new trust relationships easily. This would result in an increased volume of PII being transferred between members of the Identity Ecosystem, thus increasing privacy risk. In the police department example, the trustmark relying party may share investigatory information more often once it receives assurances through the trustmark framework that the other police departments have the proper trustmarks for interaction. This should not be understood as a privacy problem, however. Instead, it is precisely the goal of the NSTIC to foster greater sharing and inter-operability where proper trust has been achieved. Moreover, sharing within a trustmark framework will be substantially more secure than sharing under some environments in today's status quo, especially those environments with implicit trust frameworks. The trustmark framework will create a series of

accountability mechanisms that may not currently exist in these environments, so the risk of data breach should be lower than under the status quo.

Lastly, individuals can fill the role of trustmark relying parties. The next section discusses this case in more detail.

2.3 Individuals as Trustmark Relying Parties

Recall that a trustmark relying party specifies a set of trustmarks that a trustmark recipient must have in order for the trustmark relying party to trust the trustmark recipient. For example, an individual may require that an attribute provider possesses a certain set of security trustmarks prior to signing up to use the provider's services. As another example, before providing any personal information to a Web service, an individual could require that the service adheres to a certain privacy policy, which would be indicated by a collection of privacy trustmarks.

The trust requirements that an individual has could be collected and maintained by the individual.[1] In most use cases, we envision that these requirements will be stored on a user's local devices or in third-party storage accounts controlled by the user. In typical cases, the requirements will be used only by browsers or other local applications that simply verify the requirements are met when interacting with different sites, and warn the user whenever mismatches with the requirements are detected. We do not envision a need for individuals to openly publish their trust requirements nor for these requirements to be reviewed by identity providers, attribute providers, and relying parties. With typical use cases, there is no privacy risk involved with an individual maintaining their trust requirements for use within the trustmark framework.

If, for some reason, individuals want to publish their trust requirements, they may be revealing certain information, such as their name and contact information. We reiterate that publication of these requirements by an individual is a completely voluntary action and there is no requirement for publication in order for an individual to participate in the trustmark framework.

3 Privacy Benefits in the Trustmark Framework

Along with zero or minimal privacy risk in the trustmark framework, there are potential privacy benefits. As mentioned earlier, information sharing within trust frameworks supported by trustmarks will be significantly more secure than ad hoc sharing arrangements that do not have rigorous trust support. Also, adoption of the trustmark framework will lead to more identity reuse and less proliferation of user accounts, which will reduce the exposure to possible privacy violations.

Many of the initial trustmarks in the GTRI project will be on security and interoperability topics, such as levels of e-authentication assurance [11], single-sign on profiles [12], and minimum security requirements for information systems [13]. However, we are also exploring the possibility of developing privacy

[1] In the GTRI Trustmark Framework, these requirements may be maintained in a Trust Interoperability Profile (TIP).

trustmarks. Deployment of privacy trustmarks has the potential to provide a much higher level of privacy assurance to users than they receive via current practices such as an organization simply posting its privacy policy on its Web site.

3.1 Overview of Privacy Principles

Concerns over protection of personal information that is stored electronically date to the early days of computers [3]. These concerns led to various statements that came to be known as "Fair Information Practices", which entities holding personal information about individuals were suggested, and in some cases required, to follow. The first set of fair information practice (FIP) principles was codified in the Privacy Act of 1974 and set the requirements for all U.S. federal government agencies. Although these principles were left unnamed, basic notions of transparency, accuracy, consent, and security were included, laying the foundation for numerous follow-ons. Major broad-based FIP statements that are currently relevant include:

- the Organization for Economic Cooperation and Development (OECD) Privacy Principles (http://oecdprivacy.org/),
- the Consumer Privacy Bill of Rights (CPBR) [4] proposed by the U.S. executive branch,
- the U.S. Department of Homeland Security (DHS) FIP Principles [5],
- the HIPAA Privacy Rule for U.S. health care and health insurance providers [6],
- the Asia-Pacific Economic Cooperation (APEC) Privacy Framework (http://www.ag.gov.au/RightsAndProtections/Privacy/Documents/APECPrivacyFramework.pdf), and
- the Association for Computing Machinery (ACM) Privacy Recommendations (http://usacm.acm.org/privsec/category.cfm?cat=7).

Additional statements that are specifically relevant to the NSTIC Program include:

- the NSTIC FIP Principles (http://www.nist.gov/nstic/NSTIC-FIPPs.pdf), which are nearly identical to the DHS FIP Principles,
- a set of privacy requirements derived by the NSTIC National Program Office from the NSTIC strategy document (http://www.idecosystem.org/filedepot_download/28/851),
- the Federal Identity, Credential, and Access Management (FICAM) Privacy Policy [7], and
- a set of criteria proposed for privacy evaluation (http://www.idecosystem.org/filedepot_download/1433/1090) by the Privacy Coordination Committee of the Identity Ecosystem Steering Group (IDESG).

Comparing the high-level principles proposed in these different statements shows that there is quite a bit of overlap but not total consensus on fair information practices.

OECD's eight FIP principles are:

- Collection Limitation,
- Data Quality,
- Purpose Specification,
- Use Limitation,
- Security Safeguards,
- Openness,
- Individual Participation, and
- Accountability.

CPBR includes:

- Individual Control,
- Transparency,
- Respect for Context,
- Security,
- Access and Accuracy,
- Focused Collection, and
- Accountability.

The DHS FIP Principles are:

- Transparency,
- Individual Participation,
- Purpose Specification,
- Data Minimization,
- Use Limitation,
- Data Quality and Integrity,
- Security, and
- Accountability and Auditing.

The HIPAA Privacy Rule includes:

- Individual Access,
- Correction,
- Openness and Transparency,
- Individual Choice,
- Collection, Use, and Disclosure Limitation,
- Safeguards, and
- Accountability.

The APEC Privacy Framework consists of:

- Preventing Harm,
- Notice,
- Collection Limitation,

- Uses of Personal Information,
- Choice,
- Integrity of Personal Information,
- Security Safeguards,
- Access and Correction, and
- Accountability.

The ACM Privacy Recommendations include the following categories:

- Minimization,
- Consent,
- Openness,
- Access,
- Accuracy,
- Security, and
- Accountability.

The NSTIC Derived Requirements includes the following topics:

- Limiting the collection, use, aggregation, and retention of PII,
- End user notice,
- Data access and accuracy,
- Addressing and honoring end users' choices and grievances,
- Accountability,
- Supporting anonymous and pseudonymous identities, and
- Voluntary participation.

Finally, the FICAM Privacy Policy contains:

- Opt-in,
- Minimalism,
- Activity Tracking,
- Adequate Notice,
- Non Compulsory, and
- Termination.

The IDESG privacy evaluation criteria are significantly different than the privacy principles from these other organizations, primarily because they are intended for a wholly different purpose. The main purpose of the IDESG criteria is to facilitate the evaluation of privacy risks and associated mitigations for Identity Ecosystem participants, whereas the privacy principles are more general principles that are intended to be applied to all aspects of an organization's handling of personal information. Due to their focus on evaluation, the IDESG criteria partition operations into the different lifecycle phases of information, i.e. collection, use, disclosure, and retention. The IDESG criteria then look at different actors and relationships, types of information, information uses, data flows, and legal and regulatory requirements for each of the different phases. Since these criteria are intended to drive privacy evaluation rather than to provide a prescriptive solution, we do not foresee specific trustmarks being developed for

these criteria. However, the evaluation approach described by IDESG can certainly provide guidance as to the assessment criteria for different trustmarks, an aspect that must be specified within each trustmark definition.

3.2 Privacy Trustmark Granularity and Composability

We have seen in Section 3.1 that a variety of different privacy frameworks are used by different communities and in different contexts. One of the primary goals of a trustmark framework is to allow standard trust components to be used across communities and within different contexts. Given the variety of privacy frameworks in use today, it is not immediately obvious whether this can be accomplished and, if so, how. In this section, we perform a preliminary analysis of portions of the privacy frameworks discussed in Section 3.1. This analysis is intended to explore whether there are core underlying privacy criteria that these different frameworks have in common, which could form the basis for a set of interoperable privacy trustmarks. Since there is no consensus on the highest-level privacy principles included in these frameworks, the core privacy criteria will, by necessity, be finer-grained than the high-level privacy categories discussed earlier. For privacy trustmarks to be viable as a partial solution to the problem of scaling trust, one should be able to aggregate these fine-grained privacy trustmarks in trust interoperability profiles that would ensure compliance with the different privacy frameworks.

Our analysis in this paper focuses on the aspect of minimization, which is an important principle, somewhat akin to the principle of least privilege in security systems [8]. Generally speaking, minimization refers to minimizing the data that are collected and stored, and limiting uses of data to as small a set as possible while still fulfilling the intended functions of a system or service. Adhering to minimization can be seen as reducing risk by reducing the opportunities for data to be inappropriately disclosed and limiting the damage that results from inappropriate disclosures. The next paragraphs present a detailed analysis of the minimization aspects of all of the privacy frameworks discussed in Section 3.1.

OECD addresses minimization in its "Collection Limitation" principle but only states that "There should be limits to the collection of personal data". Under the "Purpose Limitation" principle, OECD states "The purposes for which personal data are collected should be specified not later than at the time of data collection and the subsequent use limited to the fulfillment of those purposes". Finally, under the "Use Limitation" principle, "data should not be disclosed … for purposes other than those specified" with special exceptions if the consent of the user is explicitly obtained or when required by law enforcement.

CPBR's "Respect for Context" principle limits both use and disclosure of data by companies to "purposes that are consistent with the relationship that they have with consumers and the context in which consumers originally disclosed the data". The "Focused Collection" principle adds that data collection should be limited to the same purposes as mentioned in "Respect for Context", while adding a specific exception for law enforcement. CPBR's "Focused Collection" adds an additional stipulation that data should be securely disposed of or de-identified once it is no longer needed.

DHS FIP principles state that only data required for the specified purpose should be collected and that use and disclosure (sharing) should be limited to that purpose or compatible purposes. DHS also specifies that data should be retained only as long as it is necessary for the stated purpose.

The HIPAA Privacy Rule includes the principle of "Collection, Use, and Disclosure Limitation". Because of its specific context, HIPAA deals only with protected health information, or PHI. The aforementioned principle states that collection, use, and disclosure of PHI should be minimized and it details specific uses and disclosures that are allowed under HIPAA, stipulating that uses/disclosures beyond those specified require explicit user (patient) consent.

APEC principles are fairly closely aligned with OECD but require that data collection (in addition to use) be limited to the specified purpose. Recall that OECD only stated that "there should be limits" on data collection.

ACM's Privacy Recommendations state that data should be collected and used only for purposes explicitly stated in an organization's privacy policy, and that data should be removed when it is no longer needed for those purposes. ACM goes beyond other privacy frameworks in several ways: it states that organizations should "implement mechanisms to evaluate, reduce, and destroy unneeded and stale info" and they should also "evaluate new activities and technologies for effectiveness, necessity, and proportionality".

The NSTIC Derived Requirements include statements on minimization that essentially mirror those in the DHS FIP Principles. In addition, the NSTIC Derived Requirements call for the minimization of data aggregation, data linkages across transactions, and data retention.

FICAM's Privacy Policy deals primarily with the context of federated authentication and so the data considered are primarily attributes used for authentication and authorization within a federated environment. Its rules state that authenticating entities can only request attributes for those purposes (authentication and authorization) and that attribute-providing entities can only disclose attributes that are requested for those purposes. Use of the attributes internally within the various entities is also limited to those purposes (a related policy employed by the National Identity Exchange Federation (NIEF) extends valid uses to include auditing as well).

Careful examination of these various minimization criteria reveals that, while specific purposes of data collection, use, and disclosure[2] differ depending on the context, there are a number of commonalities in the types of limitations included. In fact, the following small set of principles cover almost all of the minimization aspects included in the considered frameworks:

1) *minimize collection*: collect only data needed for stated purposes
2) *limit use*: use data only for stated purposes
3) *minimize disclosure/sharing*: disclose and/or share only data required for a given transaction
4) *limit disclosure/sharing*: disclose/share data only for stated purposes
5) *minimize lifetime*: delete or de-identify data that is no longer needed for stated purposes

[2] Note that use and disclosure differ in that uses can be internal to an organization without ever disclosing data. Use can also involve disclosing aggregated data without disclosure of individual data items. Thus, use is a more general concept than disclosure.

It is our belief that, in the area of minimization, interoperable privacy trustmarks covering these 5 aspects could be developed and they would cover the vast majority of minimization-related requirements from different privacy contexts. A few custom trustmarks would still be required to handle less-common minimization aspects, e.g. evaluation of activities and technologies for necessity and proportionality.

Another aspect that could affect trustmark composability is the level of assessment rigor required by different trustmarks. If composing trustmarks having different levels of assessment rigor is not viewed as a problem (the level of assessment rigor required is spelled out in the individual trustmark definitions), then this is a non-issue. However, in certain cases, combining trustmarks with different assessment standards could be undesirable, which would pose another barrier to trustmark composition. This issue is outside of the scope of this paper but will have to be dealt with when composing trustmarks in practice.

4 Conclusions

Many approaches to federated identity management create privacy risks, and it is thus important to analyze the privacy risks and benefits of new Identity Ecosystem concepts. This paper has shown why the GTRI trustmark framework has very low privacy risks, primarily because the information exchanged through the framework overwhelmingly concerns organizations, rather than individuals' personally identifiable information. At the same time, the trustmark framework can provide significant privacy benefits, to the extent that the preliminary work on creating privacy trustmarks matures into workable privacy trustmarks in practice.

While the privacy requirements analysis herein covered only the minimization aspects of privacy, it is our belief that similar analyses can be done for other privacy aspects and that the basic conclusion will hold, namely that a fairly small set of common privacy trustmarks can be developed that will cover the vast majority of privacy requirements in different contexts. However, a small number of custom trustmarks dealing with less common privacy requirements might be necessary in some cases for complete coverage. Given the large number of common privacy requirements, it should be possible to define trustmarks covering these, which will then be usable across multiple contexts and communities within the Identity Ecosystem. Some benefits of this would be standardization of privacy policy specification, interoperability that would benefit automated trust negotiation systems [9], and increased clarity for users on different privacy policies and the levels of privacy provided by different entities with whom they interact.

References

[1] GTRI NSTIC Trustmark Pilot Web Page, https://trustmark.gtri.gatech.edu/.

[2] Peter Swire and Cassandra Butts, "The ID Divide: Addressing the Challenges of Identification and Authentication in American Society," Center for American Progress (2008), http://www.americanprogress.org/issues/civil-liberties/report/2008/06/02/4520/the-id-divide/.

[3] R. Gellman, "Fair Information Practices: A Basic History," v2.11, April 4, 2014 (available at http://www.bobgellman.com/rg-docs/rg-FIPShistory.pdf).

[4] Consumer Data Privacy in a Networked World: A Framework for Protecting Privacy and Promoting Innovation in the Global Digital Economy, http://www.whitehouse.gov/sites/default/files/privacy-final.pdf.

[5] Privacy Policy Guidance Memorandum, Memorandum Number 2008-01, Department of Homeland Security, http://www.dhs.gov/xlibrary/assets/privacy/privacy_policyguide_2008-01.pdf.

[6] Department of Health and Human Services, Final Rule, Standards for Privacy of Individually Identifiable Health Information, 65 Federal Register 82462, 82464 (Dec. 28, 2000). (available at http://www.gpo.gov/fdsys/pkg/FR-2000-12- 28/pdf/00-32678.pdf)

[7] FICAM Trust Framework Solutions: Trust Framework Provider Adoption Process for All Levels of Assurance, v2.0.2, March 14, 2014.

[8] Jerry H. Saltzer and Mike D. Schroeder, "The protection of information in computer systems," *Proceedings of the IEEE,* **63** (9): 1278–1308, Sept. 1975.

[9] F. Paci, D. Bauer, E. Bertino, D. Blough, A. Squicciarini, and A. Gupta, "Minimal Credential Disclosure in Trust Negotiations," *Identity in the Information Society*, pp. 221-239, October 2009.

[10] IETF RFC 5280: Internet X.509 Public Key Infrastructure Certificate and Certificate Revocation List (CRL) Profile, May, 2008.

[11] NIST Special Publication 800-63-2, Electronic Authentication Guideline, August 2013.

[12] Profiles for the OASIS Security Assertion Markup Language (SAML) V2.0, OASIS Standard, March 15, 2005.

[13] NIST Federal Information Processing Standards Publication 200, Minimum Security Requirements for Federal Information and Information Systems, March 2006.