Enabling Scalable Multidimensional Trust in Heterogeneous Distributed Systems with Machine-Readable Trustmarks

Matthew Moyer John Wandelt Stephen Goldstein Jeff Krug
Brad Lee Stefan Roth Drew Taylor E. Anwar Reddick

Georgia Tech Research Institute
Atlanta, GA, USA

March 15, 2016

Abstract

Thanks to the innovative work of security researchers and standards bodies, it is now possible to implement a wide variety of secure transactions and use cases online using standard protocols and technologies. System implementers can now employ security features such as strong data encryption, digital signatures, secure communication channels, federated identity assertions, and permission authorizations and delegations using well-vetted open standards and open source software libraries. But these technologies are useful only in the context of appropriate trust relationships, and we contend that the current state of practice in designing and implementing Internet trust frameworks is inadequate. As the Internet continues to grow and support an increasingly large and diverse set of secure transactions involving increasingly diverse sets of trusted partners, existing trust paradigms will be unable to support the richness, scale, and heterogeneity of trust decisions that individuals and organizations will need to make. We have identified a need for a robust and highly scalable trust framework that can enable a diverse set of individuals and organizations to make fine-grained trust decisions. To meet this need, we have developed the concept of a *trustmark*. Abstractly, a trustmark is a machine-readable, cryptographically signed digital artifact representing that a specific named entity conforms to a well-scoped set of requirements, as attested

by a trusted third-party assessor. Building on this concept, we have designed, implemented, and piloted an agile and scalable *trustmark framework* that enables organizations and individuals to make fine-grained trust decisions about each other for the purpose of engaging in online transactions. In this paper, we introduce the trustmark framework, describe how it helps to solve problems related to scalable trust, and summarize our experiences with it and our roadmap for future work in this area.

1 Introduction

The Internet is in a continual state of growth and maturation, adapting in scale and scope to enable an ever-increasing set of use cases that involve trusted communication, collaboration, and data exchange. The past 20 years have witnessed substantial progress in the development of fundamental standards and technologies needed to enable these use cases. Such technologies include TLS ([1]) for secure communication channels, the XML and JSON standards ([2],[3]) for transmitting structured data payloads, XML and JSON encryption and signing standards ([4],[5],[6],[7]) for protecting the confidentiality and integrity of that data, SAML ([8]) and OpenID Connect ([9]) for federating identities across systems and applications, and OAuth 2 ([10]) and UMA ([11]) for enabling delegation of authorization and consent to systems and resources under various circumstances, plus the lower-level cryptographic primitives upon which these technologies rely. Despite their capabilities, these technologies represent only part of the solution to enabling trusted online use cases, because each of them implicitly relies on the existence of a *trust framework* to ensure that transactions can happen only among parties that have chosen to trust each other. To date, a handful of trust framework paradigms have emerged, including public key infrastructure (PKI) ([12]), bilateral (pairwise) trust arrangements, and multi-party "federated" trust frameworks. But we contend that none of these existing paradigms is sufficiently agile to capture the richness and nuance of real-world trust relationships while also scaling to accommodate the complex web of interrelationships that exists among individuals and organizations. To address this shortcoming, we have designed and implemented a solution — a robust framework based on the concept of a *trustmark* — and we are piloting the framework in the context of trusted information exchanges within a segment of the U.S. law enforcement community. Our results to date are promising, and we plan to continue developing and expanding the scope of the framework to serve other communities in the near future.

1.1 The Trustmark Concept in Brief

We define a *trustmark* as a machine-readable, cryptographically signed digital artifact representing that a specific named entity conforms to a well-scoped set of requirements. A trustmark is a form of certification, obtained in a well-defined and trusted manner that allows it to be used with confidence by other entities that may rely on it. The value of a trustmark is that it represents a trusted third-party attestation about certain salient trust characteristics of an entity, such as the policies, procedures, and technologies that it upholds or implements. A trustmark is similar to a PKI certificate, in that it is an official statement made about the trustworthiness of a specific entity, with regards to a specific set of evaluation criteria, made as a result of a thorough assessment of that entity against the evaluation criteria. But whereas a PKI certificate represents a limited and static set of facts about its subject, a trustmark can be defined based on any characteristic or set of characteristics that an entity may possess, as long as the characteristics can be clearly expressed and verified by a third party. There can be many types of trustmarks — we have defined over 600 different types so far through a pilot project — and the set of trustmarks held by an entity represents a set of third-party-verified characteristics of the entity, which other entities can use to make decisions about whether and how much to trust the trustmark holder.

1.2 Notable Contributions and Characteristics of Our Work

Our research builds upon existing trust framework paradigms, including PKI, bilateral trust agreements, and multi-party trust federations, to create a more robust, agile, and scalable framework for making trust decisions in support of online transactions. Notable contributions and characteristics of our work include the following.

1. We have developed a *trustmark framework*, including an abstract model and a normative technical specification that enables the model to be implemented in practice. We have also developed a *trustmark legal framework*, comprising a *trustmark policy* and accompanying legal agreements, to enable organizations to implement the abstract model in a real-world legal and business context.
2. Our framework enables the *componentization of trust characteristics and requirements* at an arbitrary level of granularity, which encourages wide-scale reuse of trust components and arbitrary recomposition of trust components as needed to support various business and technical use cases. It

also enables entities to publish their trust requirements explicitly and unambiguously, in a transparent manner, as a profile or collection of trust components that potential trusted partners must exhibit.

3. Our framework enables trust components to be formalized in a machine-readable format as *trustmark definitions*. In addition, our framework enables *trustmarks* — statements of conformance to specific trustmark definitions — to be published in a machine-readable format. This machine readability enables automated reasoning about the precise difference between one entity's trust requirements and another entity's capability to fulfill those requirements. We use the term *residual risk* to describe this difference, as it provides a meaningful measure of the remaining risk that an entity must accept, should it choose to participate in transactions with a specific partner entity, after accounting for all applicable trustmarks held by that partner entity.
4. Our framework requires that each trustmark definition specify a set of *conformance criteria*, which formally define the trustmark's scope of trust requirements. It also requires each trustmark definition to specify a set of formal *assessment steps* that an assessor must complete, along with specific evidentiary artifacts that the assessor must collect and maintain during the assessment, before the assessor can claim that an entity qualifies for the trustmark. These requirements encourage properties of *equivalence* and *fungibility* among trustmarks of any given type, so that, for example, a trustmark of type X, issued by Assessor A, is functionally equivalent to a trustmark of type X, issued by Assessor B.
5. By enabling componentization of trust requirements and issuance of machine-readable trustmarks corresponding to those trust components, our framework enables the creation of fine-grained trust components, where appropriate, which leads to a high degree of trustmark *reusability*. For example, suppose an organization obtains a trustmark pertaining to its physical security policies and practices, for the purpose of engaging in Use Case X with Trusted Partner A. The organization can then subsequently reuse the same trustmark, if appropriate, for the purpose of engaging in Use Case Y with Trusted Partner B. This property of reusability can lead to substantial cost savings by enabling holders of trustmarks to leverage them across a wide variety of use cases.
6. Leveraging our framework, we have developed and begun to implement a long-term vision for a distributed, scalable *trustmark marketplace* in which a wide variety of participants can acquire trustmarks and rely upon the trustmarks of other participants for a diverse array of use cases that span

communities, trust frameworks, and technologies. Also, through the course of a trustmark pilot project within the U.S. law enforcement community (see [13]), we have begun to seed this marketplace and develop the necessary infrastructure — including over 600 trustmark definitions covering basic principles of security, privacy, and identity assurance — to support its growth and maturation.

We contend that the combination of these characteristics makes our work a substantial and novel contribution to the research literature in trusted distributed systems.

1.3 Structure of This Paper

The remainder of this paper proceeds as follows. First, Section 2 describes the inherent challenges that we seek to overcome. Section 3 then discusses the shortcomings of existing approaches to trust establishment, and Section 4 highlights previous research related to our work. Next, Section 5 presents our solution — the trustmark framework — in detail, Section 6 summarizes our experiences to date with the trustmark framework, and Section 7 explains how trustmarks can be used to overcome the limitations of prior approaches in addressing the challenges. We then conclude the paper with a description of our long-term vision for a trustmark marketplace in Section 8, a summary of our future plans in Section 9, and our final remarks in Section 10.

2 The Challenge of Enabling Scalable Multidimensional Trust

The Internet is a heterogeneous distributed system, and it is becoming increasingly complex and interconnected with each passing day. People and organizations need to interconnect, communicate, and share information via the Internet in increasingly sophisticated use cases. Accordingly, over time we have seen the development of a variety of new *enabling technologies* — standards, protocols, etc. — to support these use cases. Within the context of these use cases, trust requirements and concerns — in particular, those that relate to identity, security, and privacy — are now nearly universal, regardless of the use case, and a rich ecosystem of enabling technologies exists to provide technical solutions that address these concerns. We have TLS for session confidentiality and mutual client/server authentication, SAML and OpenID Connect for identity assurance, XML Encryption and JSON Web Encryption for data confidentiality, XML Signature and JSON Web Signature for data integrity, OAuth 2 for delegation of authorization, and UMA for delegation of consent and privacy, to name only a few examples. Some of these specifications are already well-established in the marketplace, and others are just

emerging, but each one seeks to provide a standard, well-accepted technical mechanism to enable implementers to satisfy the specific requirements of their target use cases.

Across all these use cases, there exists a need for appropriate levels of trust among the use case participants. Trust is the fundamental foundation of all other properties that we expect our transactions to respect. Promises about security, privacy, identity, and other properties are meaningless without adequate trust in the entity that is making the promises. Unfortunately, as a rule, enabling technologies tend to assume the existence of an a priori trusted relationship between the participating entities as a precondition for their proper application. Such trusted relationships are straightforward to establish, as long as the criteria for establishing the relationship are simple and homogeneous and the number of trust relationships is small. But neither of these conditions is true in today's world. The digital online “trust landscape” is a richly interconnected web of trust relationships, and almost no organization or community on the Internet is an island unto itself. Consider, for example, the complex set of interrelationships among government agencies within the United States. Law enforcement agencies at the federal, state, and local levels must collaborate and share data with each other, as well as with public safety and first responder organizations. Public safety and first responder organizations must collaborate and share data with agencies and private companies that are involved in critical infrastructure protection, e.g., electric power, oil and gas, etc. Agencies at the federal level must collaborate and share data with state and local agencies. State and local agencies must collaborate and share data with other state and local agencies across state and local jurisdictional boundaries. Many government agencies must collaborate and share data with health care organizations. The interrelationships among businesses in the private sector are similarly complex, and all of these organizations, public and private, must also engage in online transactions with individual citizens and consumers. To enable trust in such a rich and complex environment, we require a framework or solution that is open, standards-based, scalable, and agile enough to satisfy a diverse set of real-world trust requirements across a wide range of participants, communities, technologies, and use cases.

3 Shortcomings of Existing Approaches to Trust Establishment

There are currently three widely used approaches to establishing the necessary trust to underpin online digital communications. These are *public key infrastructure* (PKI), *bilateral trust agreements*, and *federated trust frameworks*. This section discusses the shortcomings of each of these approaches.

PKI The primary shortcoming of PKI as a trust model is that it is not rich enough or agile enough to permit meaningful trust decisions in a multidimensional trust context. Commercial PKI certificate authorities (e.g., VeriSign, GoDaddy, etc.) provide a valuable service by issuing certificates to websites and enabling encrypted traffic for e-commerce and other use cases. But their certificates offer little more than a guarantee that the holder of the certificate is a legitimate business or individual that owns a DNS domain name. What if, for example, we want to know whether the operator of a website adheres to a sensible set of IT security principles or upholds a specific privacy practice? PKI-based trust solutions cannot address this concern, because PKI technology is inherently one-dimensional and binary: an entity either has a PKI certificate or does not have one.

Bilateral Trust Agreements As the Internet continues to be used for an ever-increasing variety of business use cases requiring trust, we are seeing the proliferation of *ad hoc* trust relationships between entities. The most common such relationship is a bilateral arrangement in which two entities (e.g., two businesses or two government agencies) agree to trust each other for a specific set of digital interactions. These relationships typically encompass multiple dimensions of trust concerns, including security, privacy, etc. The participants may execute a legal agreement to underpin the trust, or they may rely on pre-existing informal trust or goodwill. Such relationships often require significant time to establish, as they tend to require custom agreement language and approval by attorneys and business executives. In general, an ad hoc approach to building trust relationships is slow, unscalable, and non-agile.

Federated Trust Frameworks To circumvent the scaling problems inherent in a bilateral trust strategy, some communities have begun to implement multi-party *federated trust frameworks*, in which each participating entity agrees to implement a common set of policies, procedures, and technologies. Typically, a trust framework is operated by a business entity — the *framework manager* or *framework operator* — and each framework participant executes a legal agreement with the operator. Examples of multi-party trust frameworks include the InCommon

Federation ([14]), which serves the U.S. higher education and research community, and the SAFE-BioPharma Association ([15]), which serves the pharmaceutical industry. Federated trust frameworks are clearly more scalable than pairwise trust agreements, as they enable multidimensional N×N trust among participating entities with only one legal agreement required per participant. But their scalability comes at the cost of agility, as each participant must adhere to a rigid set of rules that may not meet its full range of requirements across all its business use cases. The problem of homogeneity is limiting enough when a federated trust framework serves just one community or sector, but it gets even worse when organizations in a federated trust framework need to participate in trusted transactions that span multiple communities. Consider, for example, a university that wants to participate in both the education community and the pharmaceutical research community. To do so through a federated trust framework approach, it could participate in both InCommon and SAFE-BioPharma. But these are two wholly different and incompatible trust frameworks, so the university would need to take steps to cleanly segregate its interactions within each framework to avoid violating the requirements stipulated by either framework.

4 Related Research

In addition to the existing trust establishment approaches discussed in Section 3, there are also several research efforts related to establishing online trust. We survey such related work in this section.

Blaze et al. ([16],[17]) have developed the PolicyMaker and KeyNote systems, which bear several important similarities to our work. For example, they identify credentials (third-party-issued certificates or statements) as key determinants of trust in distributed systems, and they also treat trust decisions as essentially access control decisions. Both of these aspects of their work carry over into ours.

Trust negotiation ([18],[19],[20],[21],[22],[23]), in which the parties to a prospective trust relationship attempt to achieve mutual trust through the strategic disclosure of a minimal set of sensitive credentials required by their counterparties, also relates closely to our work. We do not focus on such negotiation strategies or protocols for minimizing information disclosure, because this is not a priority for the communities that we serve. (Among our target communities, most trust requirements are stipulated by law or statute, and therefore credentials representing compliance with those requirements are not considered sensitive.) But more generally, we believe that trust negotiation techniques and protocols could be integrated with our work in a straightforward manner, with trust credentials (trustmarks) disclosed selectively and in accordance with appropriate release policies.

In addition, we believe that our work can provide a foundational vocabulary of standard credential types and definitions for wide-scale enablement of trust negotiation capabilities and applications in the online marketplace.

The "Vectors of Trust" (VoT) project ([24]) is perhaps the most closely related research to our work. VoT defines a construct called a "Vector of Trust", which is intended to capture a number of (mostly orthogonal) dimensions pertaining to the trustworthiness of a digital federated credential (e.g., strength of identity proofing, strength of credential, strength of assertion, etc.) and express that information in a standardized syntax, using a standardized taxonomy of semantically well-defined "demarcators" and values. VoT defines a "trustmark" concept, but their concept of a trustmark differs from ours in both structure and concept. For example, in the VoT model, trustmarks pertain only to IDPs, and VoT trustmarks are also far more coarse-grained than ours. But despite these differences, they aim to solve the same basic conceptual problem that we aim to address with our framework. We are currently exploring ways to integrate the VoT work with our work.

5 Our Approach

In this section, we discuss our approach to address the challenges of scalable, multidimensional trust that we illuminated in Section 2. First, we summarize the design goals for our solution in Section 5.1. Next, in Section 5.2, we introduce the trustmark framework in detail, and in Section 5.3, we provide an overview of the trustmark legal framework that underpins the trustmark framework. Section 5.4 then points out some conceptual similarities between the trustmark framework and other well-known trust and security paradigms, and Section 5.5 briefly discusses the software tools that we have developed and are planning to develop in support of the trustmark framework. Finally, Section 5.6 describes an example "end-to-end" usage scenario for trustmarks, to illustrate how everything comes together in a working system.

5.1 Trustmark Framework Design Goals

We designed the trustmark framework to meet the following goals.

1. *Agility* - The framework must be able to meet the trust needs of a wide range of organizations and individuals from disparate communities, with a wide range of interrelationships, across a wide range of use cases.
2. *Scalability and Decentralization* - The framework must be able to grow to Internet scale, i.e., it must be able to handle thousands or millions of participating organizations, and millions or billions of participating individuals.

To meet this goal, the framework must be fully decentralized, i.e., it must not require substantial centralized coordination for its operation.

3. *Componentization of Trust Requirements* - The framework must enable the decomposition of trust requirements into arbitrary-sized components, so that users of the framework can factor trust requirements into components that best facilitate component reuse.
4. *Automatability and Machine-Readability* - The framework must enable machine-readability of various technical artifacts, including trustmarks and other objects, through appropriate technical specifications, to promote the development of software tools that can help manage the inherent complexity of the framework.
5. *Trust Requirement Bundling for Business Cases* - The framework must enable the aggregation of trust components into trust requirement bundles, or profiles, to meet the trust prerequisites of specific business cases as prescribed by organizations or individuals, or to comply with requirements prescribed by law.
6. *Implementability under a Practical Legal Framework* - The framework's design must include due consideration for the practicalities of its implementation and use within a legal context. As the framework inherently supports a large number of participating organizations and individuals, it must include mechanisms to enable those participants to enter into appropriate legal agreements underpinning the trust relationships that the framework enables.

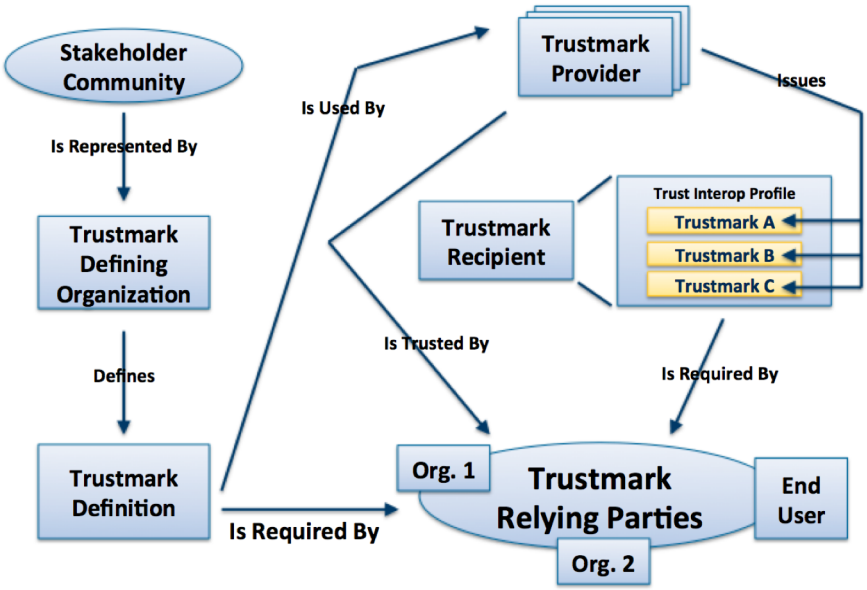

Figure 1: The Basic Trustmark Framework

5.2 The Basic Trustmark Framework

Figure 1 illustrates the basic *trustmark framework*. This section introduces the trustmark framework by defining the terms and concepts represented in the figure. Note that the framework and its concepts are formalized in a Trustmark Framework Technical Specification ([25]) that we have developed and published online. The framework comprises the following concepts and interrelationships.

Trustmark (TM) A *trustmark* (TM) is a machine-readable, cryptographically signed digital artifact that represents a statement of conformance to a well-defined set of requirements. The issuer of a TM must cryptographically sign it to ensure its integrity. [25] provides an XML-based normative specification for TM objects.

Trustmark Provider (TP) A *trustmark provider* (TP) is an entity that issues a TM based on a formal assessment process. The TM serves as a formal attestation by the TP that the recipient conforms to a well-defined set of requirements. The TM is issued under a legal framework, which we discuss in Section 5.3. Any number of TPs can exist in the framework.

Trustmark Recipient (TR) A *trustmark recipient* (TR) is an entity that receives a TM from a TP. Note that a TR is always an organization or other business entity; TMs are not intended for issuance to individuals.[1]

Trustmark Definition (TD) A *trustmark definition* (TD) specifies the conformance criteria that a prospective TR must meet, as well as the formal assessment process that a TP must perform to assess whether a prospective TR qualifies to receive a specific type of TM.[2] There can be many different types of TMs, and each type of TM has its own TD. [25] provides an XML-based normative specification for TD objects.

Trustmark Defining Organization (TDO) A TD is developed and maintained by a *trustmark defining organization* (TDO), which represents the interests of a stakeholder community. A TDO is similar in function to a standards development organization. A TDO does not play an active role in the issuance of a TM, and does not enter into any legal agreement as part of the issuance or use of TMs. Its only role is to represent a stakeholder community and publish TDs that represent the requirements of that community. Any number of TDOs can exist in the framework.

Trustmark Relying Party (TRP) Possession of a TM by a TR is required by a *trustmark relying party* (TRP), which treats the TM as formal artifact or evidence indicating that the TR meets the trust criteria set forth in the TD for the TM. When it relies on a TM, a TRP enters into a legal agreement with the TP, in accordance with the legal framework that we describe in Section 5.3. A TRP may be either an organization or an individual.

Trust Interoperability Profile (TIP) A TRP can define a *trust interoperability profile* (TIP) that expresses a trust policy in terms of a set or bundle of TMs that a TR must possess to meet its trust requirements. [25] provides an XML-based normative specification for TIP objects. Given the TIP of a TRP and a set of TMs possessed by a TR, we can compute whether a state of *trust interoperability* exists between the TRP and the TR. Note that for most real-world use cases, every participant imposes trust requirements on the other participant(s). Therefore, organizations would typically act as both a TRP (imposing requirements on other participants

---

[1] In our framework, TMs serve a purpose similar to that of attributes asserted for an individual within a federated identity assertion by an identity provider. Our framework therefore assumes that an individual does not need to be a TR, as facts about individuals can be asserted by identity providers using standard identity protocols such as SAML or OpenID Connect.

[2] As noted in Section 1.2, the inclusion of well-specified assessment processes within TDs encourages the properties of equivalence and fungibility among all TMs issued for a specific TD, so that each such TM carries the same semantic meaning regardless of which TP issued it.

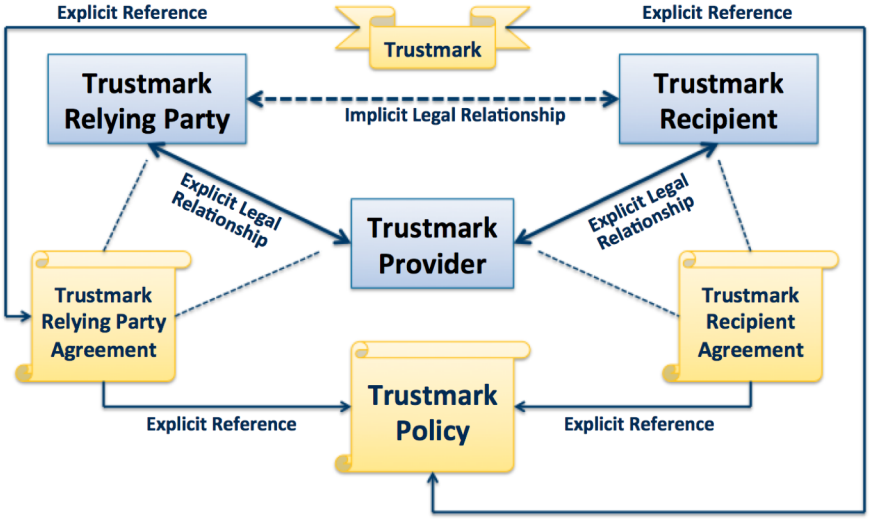

Figure 2: The Trustmark Legal Framework

through a TIP) and a TR (obtaining TMs to comply with the TIPs of other participants) within the context of a single use case.

Trustmark Status Report (TSR) After issuing a TM, a TP must publish a *trustmark status report* (TSR)[3] that provides an online, queryable source of status information about the TM. The TP must then update the TSR as needed if the TM's status changes, e.g., from "active" to "revoked" or "expired". A TRP may query the TSR periodically or as needed, based on its risk tolerance, to check whether the TM is still valid and suitable for use as a basis for trust. [25] provides an XML-based normative specification for TSR objects.

5.3 The Trustmark Legal Framework

In addition to the basic trustmark framework described in Section 5.2, we have developed a *trustmark legal framework* to support issuance, use, and reliance upon TMs as the basis of trust in live, operational use cases. Figure 2 illustrates the trustmark legal framework.

Within this framework, a TM is issued to a TR by a TP under a *trustmark recipient agreement* (TRA), which is a standard two-party contract that establishes an explicit legal agreement between the TP and TR. The TRA incorporates a *trustmark policy* by reference. The TP and TR both must sign the TRA to execute it.

[3] The TSR concept is not shown in Figure 1.

Table 1: Parallels Between the Trustmark Framework and PKI

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 Recipient Agreement	Subscriber Agreement
Trustmark Relying Party Agreement	Certificate Relying Party Agreement
Trustmark Defining Organization	N/A
Trustmark Definition	N/A
Trust Interoperability Profile	List of Trusted Certificate Authorities
Trustmark Framework Tech. Spec ([25])	X.509 Specification ([12])

When a TRP chooses to rely upon a trustmark, the TRP must enter into a *trustmark relying party agreement* (TRPA) with the TP. The TRPA is also a two-party contract; however, it is not a standard two-party agreement that both parties must sign. Instead, it is a "clickwrap" or "clickthrough" agreement that becomes effective by virtue of the TRP using or relying on a TM issued by the TP. The TRPA also incorporates the trustmark policy by reference.

Note, as indicated by Figure 2, that a TM object contains references to both the trustmark policy under which it was issued and the TRPA to which all TRPs are subject if they choose to use or rely upon the TM. Note also that even though the purpose of a TM is to provide a basis for trust between the TR and TRP, the trustmark legal framework does not establish an explicit legal relationship between these two entities. Instead, it establishes separate explicit legal relationships between each entity and a third party, the TP, which issued the TM.

Establishment of a suitable trustmark policy, TRA, and TRPA are mandatory for the issuance of a TM, as stipulated in [25]; however, [25] does not provide any further requirements or guidance as to what structure these three documents must follow or what content they must contain. We have developed versions of these legal framework artifacts that are suitable for our trustmark pilot project within the context of Georgia Tech's status as a U.S. state university, and to date we have executed 11 TRAs with various agencies from the U.S. law enforcement community.

5.4 Conceptual Parallels with PKI and ABAC

There are strong conceptual parallels between the trustmark framework and the public key infrastructure (PKI) framework, as well as between the trustmark framework and the attribute-based access control (ABAC) paradigm. We briefly explore these parallels in this section.

Table 1 illustrates how various concepts from the trustmark framework map to similar concepts from the PKI framework. There are several obvious conceptual similarities between the trustmark framework concept and the PKI framework, including the following.

1. A TM (certificate) represents a specific set of facts asserted to a TRP (certificate relying party, or audience) about a TR (subscriber) by a TP (certificate authority).
2. The roles, responsibilities, and terms of use for a TM (certificate) are described in a trustmark policy (certificate policy).
3. The scope and terms of the legal agreement between the TP (certificate authority) and the TR (subscriber) are delineated in a TRA (subscriber agreement).
4. The scope and terms of the legal agreement between the TP (certificate authority) and the TRP (certificate relying party, or audience) are delineated in a TRPA (certificate relying party agreement).

As indicated in the last few rows of Table 1, there are also other similarities between the trustmark framework and the PKI model. For example, the TFTS ([25]) maps to RFC 5280 ([12]), which defines the latest version of the X.509 specification. Also, the TIP concept loosely maps to the PKI concept of a list of trusted certificate authorities, as would be specified within a Web browser. Note that the TDO and TD concepts have no parallel in the PKI framework, as PKI is inherently one-dimensional and TDs are explicitly intended to support many dimensions of trust criteria.

Table 2: Parallels Between the Trustmark Framework and ABAC

Trustmark Framework Concept	Analogous Concept from ABAC
Trustmark	Attribute
Trustmark Definition	Attribute Definition
Trustmark Provider	Identity Provider or Attribute Provider
Trustmark Recipient	Subject or User
Trustmark Relying Party	Service Provider or Relying Party
Trust Interoperability Profile	Access Control Policy
Trust Decision	Access Decision

The trustmark framework also has strong parallels to ABAC, which is a popular and well-accepted paradigm for managing and enforcing access control policies. Table 2 illustrates them. Conceptual similarities between the trustmark framework concept and the ABAC paradigm include the following.

1. A TM (attribute) represents a fact or claim asserted by a TP (identity provider or attribute provider) about a TR (subject or user).
2. A TRP (service provider or relying party) chooses whether to trust the information conveyed by a TM (attribute) based on its level of a priori trust in the TP (identity provider or attribute provider) that generated and conveyed the TM (attribute).
3. A TRP (service provider or relying party) uses the information conveyed by a TM (attribute) as input into a trust decision (access decision) pertaining to the TR (subject or user).
4. A TRP (service provider or relying party) can define a TIP (access control policy), which is essentially a Boolean logic statement describing which TMs (attributes) a TR (subject or user) must possess to be trusted (to be granted access).
5. It is possible, in principle, for a TRP (service provider or relying party) to rely on multiple TMs (attributes) from multiple TPs (identity providers or attribute providers) as the basis for making a trust decision (access decision) about a TR (subject or user).

5.5 Software Tools for Trustmarks

As noted in Section 5.2, we have developed normative technical specifications for many of the artifact types that exist in the trustmark framework. This permits and encourages the development of various software tools to facilitate the use of the

framework in practice. We have developed the following prototype software tools as part of a trustmark pilot project.

1. *Trustmark Definition Authoring Tool* - To aid TDOs and other TD authors in developing, publishing, and maintaining TDs, we have developed a prototype software tool that compiles structured input data into TDs conformant to [25]. With this tool, we have generated over 600 TDs and published them online ([13]). We plan to refine and augment the capabilities of this tool based on our experience with it.
2. *Trust Interoperability Profile Authoring Tool* - To aid TRPs and others that want to develop, publish, and maintain TIPs, we have developed a prototype software tool that compiles structured input data into TIPs conformant to [25]. We have used this tool to generate and publish over 200 TIPs online ([13]). As with the previous tool, we plan to refine and augment its capabilities based on our expense with it.
3. *Trustmark Assessment and Trustmark Lifecycle Management Tool* - To aid security auditing firms and others that want to play the role of TP in the trustmark framework, we have developed a prototype software tool that facilitates the process of assessing prospective TRs against one or more TDs, and then generating and cryptographically signing TMs conformant to [25] based on the results of those assessments. This tool is compatible with any TD or TIP that conforms to [25]. We have used the tool to perform several hundred TM assessments of our partner agencies as part of a pilot project, and based on these assessments, we have issued and published nearly 100 TMs to these agencies.[4] The tool also allows us to publish and manage TSRs conformant to [25] for all the TMs that we publish, so that we can revoke a TM if necessary. We are currently working to refine and augment the capabilities of this tool, based on our initial experiences with it.
4. *Trustmark Binding Registry and Binding Tools* - To enable the operational use of TMs for making trust decisions, as well as the discovery of potential trusted partners based on satisfaction of trust policies expressed within TIPs, we have developed a set of prototype TM registration and binding tools. These tools allow for the publication of trust fabric entries — endpoint descriptor data structures for live computer systems and services — and the cryptographic binding of TMs to those entities, within a publicly searchable *trustmark binding registry*. We are currently using these tools

[4] Each TM issued is subject to a trustmark recipient agreement executed between Georgia Tech and the recipient organization, as noted in Section 5.3.

as part of a pilot project within the scope of a small community of interest. Currently, our trustmark binding registry and binding tools support a limited set of trust fabric formats, including SAML identity provider and service provider entity descriptors conformant to [8]. We are continuing to refine and augment these tools based on our experiences with them.

5. *Open Source Trustmark Utilities and Libraries* - To enable the development of additional software tools that are compatible with the trustmark framework, we have designed and are planning to develop a set of utilities and libraries that simplify certain common-but-critical tasks, such as testing a TM for validity or invalidity (e.g., signature integrity, expiry check, revocation check, etc.) and testing whether a given bundle of TMs satisfies a given TIP. We plan to develop these components and release them under an open source software license, as we believe these components are important for the realization of our vision the wide-scale use of trustmarks.
6. *Firefox Browser Add-on for Trustmark-Based Online Privacy* - To conceptually demonstrate the applicability of the trustmark framework to protecting online privacy, we developed a prototype Firefox add-on module that can use TMs issued to websites to help enforce a user's privacy preferences. The module works by retrieving all TMs granted to a website, testing whether the website's TMs satisfy a user-defined TIP, and then preventing or discouraging the user from visiting the website if it does not have the requisite TMs. This module is available for download on GitHub; see [26].

As noted within this section, we intend to foster the maturation and wide uptake of the trustmark framework as a trust paradigm by continuing to refine and improve the capabilities of these software tools. Sections 8 and 9 contain more information about our long-term plans to spur the growth and adoption of the trustmark framework.

5.6 An Example Trustmark Usage Scenario

In this section, we present an example TM usage scenario to illustrate how the trustmark framework works. While this example uses actual organizations and realistic application use cases, the details presented herein about usage of the trustmark framework by these organizations are hypothetical.

Within the U.S. law enforcement community, the Texas Department of Public Safety (TX DPS) operates an online application called TXMAP, which is a multi-purpose reporting tool for geographic and geospatial data sets. TXMAP provides an interactive map of Texas, similar to Google Maps, with a variety of data layers

overlaid onto the basic map geography. TXMAP uses SAML to enable users from various agencies to access the application, and it uses ABAC to ensure that each user can see only those data layers for which he is authorized. TXMAP can provide value to multiple communities, including federal, state, and local law enforcement officers, first responders (e.g., firefighters and emergency medical technicians), members of the infrastructure protection community (e.g., the electrical power and oil and gas industries), and others. To meet its business objective of facilitating the dissemination of the right data to the right persons from each of these communities, TX DPS wants TXMAP to be available to users from a wide range of organizations across many communities. But doing so requires that TX DPS establish a trusted relationship with each such organization. The trustmark framework is ideal for facilitating trust in this scenario.

To leverage the trustmark framework, TX DPS first defines a TXMAP TIP that specifies a trust policy to which any organization must conform as a prerequisite for connecting to TXMAP. The TIP contains a list of TDs representing the types of TMs that conforming organizations must possess, in the areas of security, privacy, identity assurance, and organizational integrity (“bona fides”). The TIP also specifies a list of approved TPs that are trusted by TX DPS to issue TMs. Conforming organizations must obtain their TMs from one or more TPs on this list.

Due to the shared physical border between Texas and New Mexico, the New Mexico Department of Public Safety (NM DPS) wants to make select layers of TXMAP data available to some of its personnel. NM DPS configures a SAML identity provider (IDP) endpoint to federate its users' local identities, and then obtains a set of TMs from one or more of the TX DPS-approved TPs. The TMs obtained reflect NM DPS policies and practices for organizational security, privacy, identity assurance, etc. Each TM obtained is subject to a TRA between NM DPS and the issuing TP, and issued TMs are published by the TP, in accordance with trustmark publishing requirements stipulated in [25]. The TMs allow NM DPS to demonstrate that it conforms to the TXMAP TIP published by TX DPS.

After acquiring the necessary TMs, NM DPS arranges for its TMs to be bound to its IDP software endpoint and published in a binding registry such the one we describe briefly in Section 5.5. Binding and publishing the TMs in a registry enables data about NM DPS and its IDP software endpoint to be discovered and used not only by TXMAP, but also by other potential trusted partner organizations. The TXMAP application can now discover the NM DPS IDP endpoint in the binding registry, inspect its TMs, and determine whether the IDP satisfies the TXMAP TIP. Assuming that the NM DPS IDP does indeed satisfy the TXMAP TIP, TXMAP can now trust the NM DPS IDP, and users from NM DPS can now access TXMAP. TXMAP can

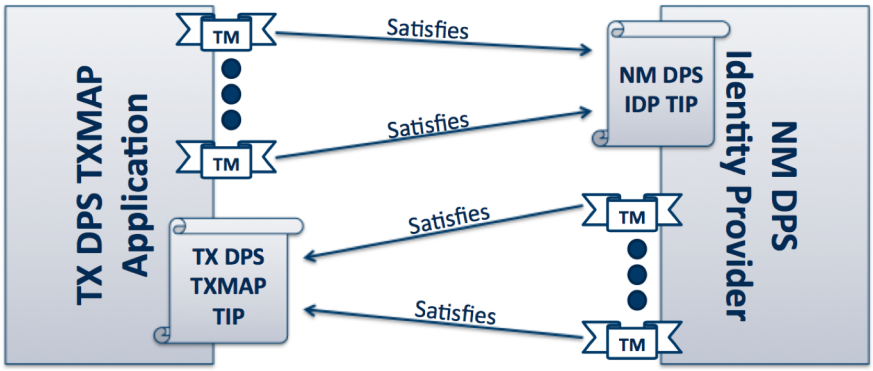

Figure 3: Mutual Trust Based on Trustmarks and TIP Satisfaction

still apply per-user, fine-grained access controls to specific TXMAP data layers, based on specific attributes conveyed about each user by the NM DPS IDP. Note that TMs could be used to facilitate both "directions" of trust in this use case. Just as NM DPS needs to satisfy the TXMAP TIP, NM DPS can also publish a TIP to which TXMAP must conform. Figure 3 illustrates this mutual TM-based trust between the two systems.

In this example, NM DPS obtains TMs to satisfy its immediate goal of enabling access to TXMAP for its users, but due to the high degree of reusability for many TDs, NM DPS is well positioned to reuse its TMs for enabling its users to gain access to other services. In the future, NM DPS may need to acquire a handful of additional TMs to enable user access to additional services, but over time, as NM DPS continues to expand its set of TM-based trust relationships with other organizations, it will likely find that it already possesses most or all of the required TMs due to trust requirements it was already required to satisfy for a pre-existing trust relationship. Over time, positive network effects can arise in which the marginal cost of establishing the (N+1)th TM-based trust relationship is small due to the high degree of reusability of work already performed to establish the first N such relationships.

6 Implementation Experience with the Trustmark Framework

Throughout Section 5, as we have discussed each aspect of the trustmark framework, we have also highlighted our specific implementation experience with it.

This section pulls together and summarizes our trustmark framework implementation experience.

1. We have executed 11 TRAs with partner agencies from the U.S. law enforcement community. These agreements allow us to issue trustmarks, and allow those trustmarks to be relied upon for operational purposes, under a real-world legal framework.
2. We have authored and published over 600 TDs, representing the componentization of a wide variety of trust criteria that are applicable to agencies within the U.S. law enforcement community. Similarly, we have authored and published over 200 TIPs that represent various logical combinations of trust criteria that correspond to specific trust and interoperability policies. These TDs and TIPs are available online at [13]. Our current set of published TDs and TIPs is based on requirements that we analyzed and derived from several policy guidance documents that are prominent within the U.S. federal government community and the U.S. law enforcement community. Specific source documents, and the high-level requirements topics that we derived from them, include the following. From [27], we extracted and requirements related to identity assurance. From [28], we extracted requirements related to identity assurance and privacy. And from [29], we extracted requirements related to organizational and system security. Specifically, we extracted only those requirements that pertain to "Low Impact" security controls.
3. We have performed several hundred TM assessments of our partner agencies. Based on these assessments, we have issued and published nearly 100 TMs, and also published binding data for most of these TMs, thereby enabling the TMs to be used with live systems and relied upon by users of those systems for real-world trust decisions.

Through our experience, we have validated the trustmark framework concept and demonstrated the value of our prototype tools in facilitating the operational use of the trustmark framework. We are continuing to gain operational experience with the trustmark framework through a pilot project, as discussed in Section 9.

7 Overcoming Limitations of Prior Approaches with Trustmarks

In Section 2, we discussed the inherent limitations of three existing approaches to establish trust for online transactions: PKI, bilateral trust agreements, and multilateral trust frameworks. In this section, we explore how the trustmark framework

can supplement each of these approaches to achieve greater scalability and diversity of trust.

PKI Our primary problem with PKI, as we have noted, is that PKI-based trust is inherently one-dimensional. This is true, for example, in the PKI infrastructure used for providing secure TLS access to websites from end-users' web browsers. We can address this problem by supplementing PKI with TMs, as follows. First, issue to each PKI subscriber (website owner) a set of TMs, in accordance with the various trust criteria to which that subscriber conforms. Next, cryptographically bind each issued TM to the website's PKI certificate, using a binding registry such as the one we discuss in Section 5.5. This allows a certificate relying party (end user with a web browser) to choose whether to trust a website not only based on its trust in the issuing certificate authority, but also based on whether the certificate has been bound to a specific set of required TMs, as expressed in the certificate relying party's TIP. This would of course require additional software capabilities, e.g., within the user's web browser. Note that while we have described this approach in the context of a specific PKI application, in principle it can work for any PKI. We can also apply this approach to peer-to-peer certificate paradigms, e.g., Pretty Good Privacy (PGP) ([30]), in which there exist no hierarchical relationships among certificates.

Bilateral Trust Agreements The primary challenge with bilateral agreements is that they scale poorly, relative to their complexity. We can address this problem by simply implementing bilateral agreements atop the trustmark framework. Using this approach, we can separate the bilateral agreement's trust requirements from its legal clauses relating to, e.g., liability, indemnification, assignment, termination, etc. We can also restructure the bilateral agreement so that its trust requirements are expressed as TIPs to which the appropriate party or parties agree to conform. Each participating organization can then acquire the necessary TMs to demonstrate that it conforms to the terms of the agreement. This approach effectively converts a verbose, complex, bilateral agreement into a concise, often boilerplate, legal agreement supplemented by a set of TIPs. The inherent machine readability and algorithmic comparability of these TIPs allows for any given organization to enter into a larger number of bilateral agreements, based on its business needs, because the time and effort required to develop each bilateral agreement and ensure conformance to it is much lower when agreements are more lightweight and their complexity is manageable by machines. In some cases, for organizations that fully leverage the trustmark framework, the inherent trustworthiness of TMs may eliminate the need for bilateral agreements entirely.

Multilateral Trust Frameworks The main problem with multilateral trust frameworks is that they tend to assume homogeneity of trust requirements and policies across all of their participants. We can address this problem by implementing these multi-party frameworks atop the trustmark framework, just as we would do for bilateral agreements. Again, by separating the trust requirements from the legal clauses, and enabling the expression of trust requirements as TIPs, we can simplify a trust framework's legal agreement(s) substantially. In addition, this approach allows for much greater flexibility in expressing and enforcing trust requirements, as each participant can articulate its own trust requirements independently from the other participants. The framework can still specify a "core" or "minimum acceptable" TIP, and participants can build upon it, defining their own custom trust requirements (TIPs) as needed in "layers" atop the core TIP. This increased flexibility with trust requirements can enable the framework to expand and serve a larger set of participants than it could otherwise serve. Also, by adopting a standardized format for expressing trust requirements, we enable any given organization to more easily participate in multiple such frameworks — bilateral or multilateral — as needed, based on its business needs. The result is a much more scalable and flexible trust environment for all participating organizations.

Note that, for both bilateral trust agreements and multilateral trust frameworks, we propose an approach by which the participants separate the legal terms and clauses of their agreement from their specific trust requirements, and express those trust requirements as trust interoperability profiles. To help facilitate this approach, our future research plans may include the development of a generic "trust agreement builder tool" that will enable organizations to author pairwise and multilateral trust agreements using the trustmark framework as the basis for expressing their specific trust requirements.

8 Long-Term Vision: A Trustmark Marketplace

Based on our early pilot results, the trustmark framework appears to be a viable concept for the enablement of agile, multidimensional trust at Internet scale. But we recognize that achieving such an ambitious goal requires more than a mere demonstration project or pilot. Our long-term vision is to catalyze the creation of a *trustmark marketplace*, based on the trustmark framework. We envision the trustmark marketplace as having the following characteristics.

1. There exists within the marketplace a robust set of TDs, covering most or all of the criteria that organizations and individuals would typically consider in making trust decisions about other organizations. Such criteria in-

clude the topics of security, privacy, identity assurance, organizational integrity ("bona fides"), and others. These criteria are componentized and published at a granularity that permits and encourages wide reuse of a "core" set of trust requirements. Beyond these core requirements, there exist additional TDs covering trust criteria that supplement the core requirements for specific communities or use cases. TDOs representing numerous stakeholder communities participate in the marketplace, developing TDs and collaborating with each other to achieve maximum reuse of the TDs that they publish.

2. The marketplace comprises a large number of organizations, representing multiple communities, interacting with and trusting each other operationally through the use of the trustmark framework.
3. Organizations participating in the marketplace can acquire TMs from a variety of third-party TPs according to a competitive market model, and each TP offers TM assessment and issuance services according to its business goals and its areas of assessment expertise. When appropriate, organizations can even make use of self-issued TMs based on a self-assessment process,[5] provided that partner organizations are willing to accept the self-issued TMs.
4. Participating communities develop and publish TIP templates that their members can leverage as a mechanism for complying with community norms and/or laws.
5. Marketplace participants are supported by a rich and thriving TM software tool ecosystem, including both open source and commercial products for tasks such as TM assessment, binding, use, reliance, revocation, etc.

Section 9 describes a series of next steps that we plan to take in the coming years, to realize this vision.

9 Next Steps

In this section, we enumerate the various threads of activity that we are currently pursuing to further develop the trustmark framework and create a thriving trustmark marketplace as described in Section 8.

Operational Pilot of the Trustmark Framework As noted in Section 6, we have already begun to pilot the trustmark framework within a select set of agencies

[5] The trustmark framework allows for self-assessment and self-issuance of trustmarks, so long as the issuing organization has assessed itself in accordance with the assessment rules specified for the trustmarks issued.

from the U.S. law enforcement community. As this pilot continues, we intend to exercise all aspects of the framework, including:

1. Identifying additional prominent policies and specifications in the areas of security, privacy, and identity assurance; decomposing requirements from those documents into generic, modular, and reusable trust components; and publishing those requirements as TDs;
2. Aggregating generic trust components into meaningful bundles based on real-world use cases and business needs, and publishing those bundles as TIPs;
3. Assessing participant agencies according to criteria specified in TDs, and issuing TMs to those agencies under actual legal agreements;
4. Binding issued TMs to endpoint descriptors representing real-world systems belonging to the agencies to which the TMs were issued; and
5. Facilitating the operational use and reliance upon TMs for making real-world trust decisions.

Our goals for this pilot are to establish the viability of the trustmark framework concept and to identify aspects of the framework that require further investigation and refinement.

Development of Additional TDs and TIPs Section 6 summarizes the scope of requirements and sources for the TDs and TIPs that we have published to date. To expand upon this initial set of requirements, we are currently developing additional TDs and TIPs derived from these additional sources.

1. From [29], we are extracting requirements related to organizational and system security for “Medium Impact” and “High Impact” security controls.
2. From [31], we are extracting requirements related to organizational and system security and identity assurance, plus additional requirements that are unique to the law enforcement community.
3. From [32], we are extracting requirements related to organizational and system security and identity assurance.

Upon completion of this work, we expect to have a reasonably complete set of generic TDs representing a rich taxonomy of trust requirements in the areas of security, privacy, and identity assurance, plus a rich set of TIPs representing trust requirements for various communities and use cases. We realize that these requirements are derived from U.S. government publications; however, based on our experience and expertise in this area, we believe that most of the TDs we are developing and publishing will be highly reusable as general-purpose trust criteria for many additional communities and use cases.

Refinement of the Trustmark Framework Technical Spec Based on our pilot project experiences to date, we have already identified a handful of refinements and improvements that we can make to the TFTS ([25]) to improve the trustmark framework. For example, [25] currently requires all TDs, trustmarks, TIPs, and TSRs to be published in XML format, but we recognize the emerging popularity of JSON as a data interchange standard, and we plan to update [25] to include support for the JSON-based suite of standards. In addition, we plan to update [25] to allow for "parameterized" trustmarks, such that a TD can represent a family of related policy rules. Parameterized trustmarks would be useful, for example, when dealing with requirements related to minimum acceptable password length. Rather than requiring a separate TD to express each possible numerical value (e.g., "Minimum Password Length of 8 Characters", "Minimum Password Length of 9 Characters", etc.), we could publish one TD and enable it to be parameterized upon trustmark issuance. We expect to publish a new version of the spec, with these refinements and others, in the near future.

Establishment of a Trustmark Initiative Due to our early outreach efforts within the communities that we serve in our work, we believe that many organizations now recognize the potential of machine-readable TMs as a fundamental enabler of scalable, peer-to-peer cross-organizational trust. But we also recognize that the impact of the trustmark framework will be limited to the scale of the community that adopts it, uses it operationally, and contributes to its ongoing refinement. Much of the inherent value in the trustmark framework concept is based on the philosophy of componentization and reuse of fine-grained trust requirements. The trustmark framework therefore works best and provides maximum value to all of its stakeholders when trust component reuse is maximized across communities and across use cases. We fear that, over the long term, without a central point of coordination, the trustmark framework stakeholder community could fracture into silos, each with its own "non-standard" TDs and TIPs. To avoid this undesired outcome, we are launching a *Trustmark Initiative*, the goal of which is long-term oversight and coordination of the following activities:

1. Ongoing development and maturation of [25];
2. Development and publication of new TD and TIP artifacts, plus harmonization of new artifacts with existing ones from across all stakeholder communities;
3. Ongoing development and maturation of software tools, such as those tools noted in Section 5.5 and others; and

4. Ongoing outreach and advocacy for the trustmark framework solution and the trustmark marketplace vision within all stakeholder communities that stand to benefit from this work.

We plan to launch the Trustmark Initiative in mid-2016. More information about it will be available at [33].

10 Conclusions

The trustmark framework provides a robust solution to the challenge of achieving scalable, multidimensional trust for a wide range of digital transactions and use cases. Unlike other prominent approaches to online trust — e.g., PKI, bilateral agreements, and multilateral trust frameworks — the trustmark framework is flexible enough to support a wide diversity of trust relationships while also remaining highly scalable. While we continue to develop the framework and pilot it with our partner agencies, our initial results indicate that the trustmark framework is viable from both a technical and a legal standpoint. We plan to further develop the framework, and facilitate its uptake and usage across a wide range of stakeholders, over the coming years.

Acknowledgment and Disclaimer

Work on our trustmark pilot, and related work that led to our trustmark pilot, was performed under awards from the U.S. Department of Commerce's National Institute of Standards and Technology (NIST), the U.S. Department of Justice's Bureau of Justice Assistance (BJA), the Office of the Program Manager for the Information Sharing Environment (PM-ISE), and the U.S. Department of Homeland Security (DHS). The statements, findings, conclusions, and recommendations are those of the authors and do not necessarily reflect the views of our sponsors.

References

[1] Internet Engineering Task Force. Request for Comments 5246: "The Transport Layer Security (TLS) Protocol Version 1.2". https://tools.ietf.org/html/rfc5246. August 2008.

[2] World Wide Web Consortium (W3C). "Extensible Markup Language (XML) 1.0 (Fifth Edition)". http://www.w3.org/TR/REC-xml/. W3C Recommendation, November 26, 2008.

[3] Ecma International. "The JSON Data Interchange Format," 1st Edition. http://www.ecma-international.org/publications/files/ECMA-ST/ECMA-404.pdf. October 2013.

[4] World Wide Web Consortium (W3C). “XML Encryption Syntax and Processing”. http://www.w3.org/TR/2002/REC-xmlenc-core-20021210/. W3C Recommendation, December 10, 2002.
[5] World Wide Web Consortium (W3C). “XML Signature Syntax and Processing (Second Edition)”. http://www.w3.org/TR/xmldsig-core/. W3C Recommendation, June 10, 2008.
[6] Internet Engineering Task Force. Request For Comments 7516: “JSON Web Encryption (JWE)”. https://tools.ietf.org/html/rfc7516. May 2015.
[7] Internet Engineering Task Force. Request For Comments 7515: “JSON Web Signature (JWS)”. https://tools.ietf.org/html/rfc7515. May 2015.
[8] OASIS Security Services Technical Committee. “Security Assertion Markup Language,” Version 2.0. http://saml.xml.org/saml-specifications. March 2005.
[9] OpenID Foundation. “OpenID Connect,” Version 1.0. http://openid.net/specs/. November 2014.
[10] Internet Engineering Task Force. Request For Comments 6749: “The OAuth 2.0 Authorization Framework”. https://tools.ietf.org/html/rfc6749. October 2012.
[11] Kantara Initiative. “User-Managed Access (UMA) Profile of OAuth 2.0,” Version 1.0.1. https://docs.kantarainitiative.org/uma/rec-uma-core.html. December 2015.
[12] Internet Engineering Task Force. Request for Comments 5280: “Internet X.509 Public Key Infrastructure Certificate and Certificate Revocation List (CRL) Profile”. https://tools.ietf.org/html/rfc5280. May 2008.
[13] Website for the Georgia Tech Research Institute's National Strategy for Trusted Identities in Cyberspace (NSTIC) Trustmark Pilot. https://trustmark.gtri.gatech.edu/.
[14] Website for the InCommon Federation. https://www.incommon.org/federation/.
[15] Website for the SAFE-BioPharma Association. http://www.safe-biopharma.org/.
[16] Blaze, M.; Feigenbaum, J.; Lacy, J., “Decentralized Trust Management,” Proceedings of the 17th IEEE Symposium on Security and Privacy. Pages 164,173, 1996.
[17] Matt Blaze, Joan Feigenbaum, John Ioannidis, and Angelos D. Keromytis. “The Role of Trust Management in Distributed Systems Security.” In Secure Internet Programming, Jan Vitek and Christian D. Jensen (Eds.). Springer-Verlag, London, UK, Pages 185-210, 2001.

[18] Elisa Bertino, Elena Ferrari, Anna Cinzia Squicciarini, “Trust-X: A Peer-to-Peer Framework for Trust Establishment,” IEEE Transactions on Knowledge and Data Engineering, Vol. 16, No. 7, Pages 827-842, July, 2004.

[19] Abhilasha Bhargav-Spantzel, Anna C. Squicciarini, Elisa Bertino, “Trust Negotiation in Identity Management,” IEEE Security & Privacy, Vol. 5, No. 2, Pages 55-63, March/April, 2007.

[20] Adam J. Lee, Marianne Winslett, Jim Basney, and Von Welch. “Traust: A Trust Negotiation-Based Authorization Service for Open Systems.” In Proceedings of the eleventh ACM symposium on Access control models and technologies (SACMAT '06). ACM, New York, NY, USA, Pages 39-48, 2006.

[21] Winsborough, W.H.; Seamons, K.E.; Jones, V.E., “Automated Trust Negotiation,” Proceedings of the DARPA Information Survivability Conference and Exposition, 2000. (DISCEX '00.) Vol.1, Pages 88-102, 2000.

[22] Marianne Winslett, Ting Yu, Kent E. Seamons, Adam Hess, Jared Jacobson, Ryan Jarvis, Bryan Smith, and Lina Yu. 2002. “Negotiating Trust on the Web.” IEEE Internet Computing Magazine, Vol. 6, Issue 6, Pages 30-37, November/December 2002.

[23] Ting Yu, Marianne Winslett, and Kent E. Seamons. “Interoperable Strategies in Automated Trust Negotiation.” In Proceedings of the 8th ACM Conference on Computer and Communications Security (CCS '01), Pierangela Samarati (Ed.). ACM, New York, NY, USA, Pages 146-155, 2001.

[24] Justin Richer and Leif Johansson. “Vectors of Trust.” Internet Engineering Task Force, Network Working Group, Internet Draft. https://www.ietf.org/id/draft-richer-vectors-of-trust-02.txt. November 2015.

[25] Georgia Tech Research Institute. “Trustmark Framework Technical Specification,” Version 1.0. https://trustmark.gtri.gatech.edu/specifications/trustmark-framework/1.0/tfts-1.0.pdf. October 2014.

[26] Trustmark-Based Online Privacy Guard Firefox Add-On Prototype. https://github.com/akshatarao/practicum.

[27] U.S. Dept. of Commerce, National Institute for Standards and Technology (NIST). “NIST Special Publication 800-63-2: Electronic Authentication Guideline”. http://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.800-63-2.pdf. August 2013.

[28] U.S. Federal Identity, Credential, and Access Management (FICAM) Trust Framework Solutions (TFS) Program. “FICAM Trust Framework Provider Adoption Process (TFPAP) for All Levels of Assurance,” Version 2.0.2.

http://www.idmanagement.gov/sites/default/files/documents/FICAM_TFS_TFPAP_0.pdf. March 2014.

[29] U.S. Dept. of Commerce, National Institute for Standards and Technology (NIST). "NIST Special Publication 800-53 Revision 4: Security and Privacy Controls for Federal Information Systems and Organizations". http://dx.doi.org/10.6028/NIST.SP.800-53r4. April 2013. Includes updates as of 15 Jan 2014.

[30] P. Zimmerman. PGP User's Guide, MIT Press, 1994.

[31] U.S. Department of Justice, Federal Bureau of Investigation, Criminal Justice Information Services (CJIS) Division. "Criminal Justice Information Services (CJIS) Security Policy," Version 5.3. https://www.fbi.gov/about-us/cjis/cjis-security-policy-resource-center. August 2014.

[32] U.S. Federal Public Key Infrastructure Policy Authority. "X.509 Certificate Policy For The Federal Bridge Certification Authority (FBCA)," Version 2.27. https://www.idmanagement.gov/sites/default/files/documents/FBCA%20Certificate%20Policy%20v2.27.pdf. December 2013.

[33] Website for the Trustmark Initiative. https://trustmarkinitiative.org/.