The Trust over IP Stack

Matthew Davie, Dan Gisolfi, Daniel Hardman, John Jordan, Darrell O'Donnell, and Drummond Reed

Abstract

This article defines a four-layer architectural stack called the ToIP stack for establishing trust between peers over the Internet and other digital networks. Patterned after the TCP/IP stack that standardized packet exchange and created the Internet, the ToIP stack is a decentralized architecture that encompasses business, legal, and technological requirements. Layer One establishes decentralized trust roots using decentralized identifiers (DIDs), an emerging W3C standard for decentralized PKI. Layer Two is the DIDComm protocol, a transport-independent protocol that uses DIDs to form and communicate over a cryptographically secure connection. Layer Three is a suite of credential exchange protocols based on the W3C Verifiable Credentials standard for cryptographically verifiable digital credentials. Layer Four adds cryptographically verifiable governance frameworks using a metamodel for describing the business, legal, and technical policies under which a peer is operating as an issuer, holder, or verifier of digital credentials. This governance metamodel can be applied at all four Layers of the stack, producing a parallel ToIP Governance Stack that fully integrates the non-technical dimensions of trust establishment. Further work on defining, testing, and integrating the ToIP stack is planned for a new project at the Linux Foundation.

Introduction

In early 2005, Kim Cameron, Architect of Identity at Microsoft, and his team observed in the *Laws of Identity* paper [1] that the Internet was created without an identity Layer. He was articulating the painful lesson that the first generation of identity and access management (IAM) professionals grappling with Internet identity had learned: this architectural shortcoming was becoming increasingly serious. In the 14 years since, the situation has only grown worse.

In a second paper released later that year [2], Cameron proposed that this Layer could be built as an *identity metasystem*:

> *Given that universal adoption of a single digital identity system or technology is unlikely ever to occur, a successful and widely employed identity solution for the Internet requires a different approach — one with the capability to connect existing and future identity systems into an identity metasystem. This metasystem, or system of systems, would leverage the strengths of its constituent identity systems, provide interoperability between them, and enable creation of a consistent and straightforward user interface to them all.*

This paper described the actors in this metasystem as having Three primary roles:

- *Identity providers* who digitally issue claims about an entity
- *Relying parties* who require access to those claims in order to do their business
- *Subjects* who are the individuals and other entities described by the claims

The paper went on to describe how such a metasystem could be implemented using two technologies that Microsoft was leading at the time:

- *The WS-* web services stack* to standardize protocols for secure information sharing
- *Information Card ("InfoCard") technology* for expressing, storing and exchanging claims-based identities using a common structured data format (XML) and a consistent user experience of selecting "cards" in an "identity selector" (much like we do with printed identity cards in our physical wallets)

Although Microsoft led a valiant charge to build this metasystem, including establishing the Information Card Foundation and developing open standards for Information Cards at OASIS, it did not ultimately succeed in the market. However, the architectural pattern advocated in these papers did not die, and now, 14 years later, an Internet identity Layer is emerging that bears a striking resemblance to that original design — with a twist.

That twist is *decentralization* — the ability for new cryptographic technologies like blockchains, distributed ledgers, distributed hash tables, and decentralized file systems to allow multiple untrusting parties to securely interact with the same universal source of truth. By using these technologies to decentralize the public key infrastructure (PKI) required for Information Cards — and by providing real market proof of the value of distributed trust — we finally have a foundation on which a durable identity metasystem can be built.

In this article we describe the new "identity stack" arising out of this architecture. We define the architectural separation between the layers, discuss how it differs from the original Information Card architecture, and explore how it is evolving.

Architectural Layering of the Trust over IP Stack

Since the ultimate purpose of an Internet identi-

Digital Object Identifier:
10.1109/MCOMSTD.001.1900029

affiliations

2471-2825/19/$25.00 © 2019 IEEE

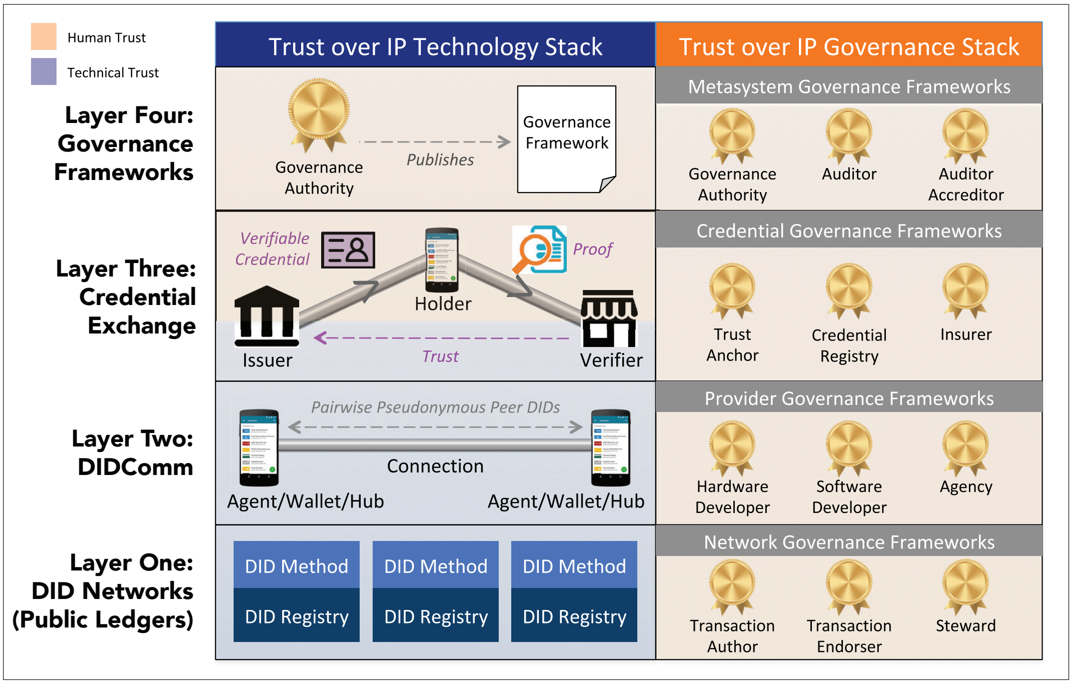

FIGURE 1. The ToIP stack.

ty layer is not actually to identify entities, but to facilitate the trust they need to interact, co-author John Jordan coined the term *Trust over IP* (ToIP) for this stack. Figure 1 is a diagram of its four Layers:

Note that it is actually two parallel stacks: technology and governance. This is a lesson learned from Information Cards — digital trust cannot be achieved by technology alone, but only by humans and technology working together. It is also a reflection that the ToIP stack evolved from the Sovrin stack defined by the Sovrin Foundation in Appendix D of the Sovrin Glossary [3]. We dig deeper into the unique role of governance in the ToIP stack when we discuss Layer Four.

Layer One: Decentralized Identifier Networks

The breakthrough in harnessing decentralization technology as the foundational Layer of the ToIP stack was the emergence of a new type of globally unique identifier called a decentralized identifier (DID). Originally developed under a research grant funded by the U.S. Department of Homeland Security Science & Technology division, since June 2017 the DID specification [4] and the DID Primer [5] have been maintained by the W3C Credentials Community Group (CCG). In September 2019 the W3C membership voted to create a DID Working Group, which then accepted the CCG specification as its starting point for producing a 1.0 standard.

DIDs combine four properties into a single RFC-3986-compliant uniform resource identity (URI) scheme:

1. *Permanence:* Once assigned to an entity (called the DID subject), a DID functions like a uniform resource name (URN), that is, it is a persistent identifier that never needs to change.
2. *Resolvability:* A DID can be resolved to a DID document describing properties of the DID subject, most notably the public key(s) and service endpoint(s) necessary to engage in trusted interactions.
3. *Cryptographic verifiability:* By registering DIDs on blockchains, distributed ledgers, or other decentralized systems and then resolving them to public keys in DID documents, a DID subject can prove cryptographic control of a DID.
4. *Decentralization:* Unlike most conventional network identifiers (e.g., phone numbers, IP addresses, domain names, email addresses), DIDs do not require centralized registration authorities.

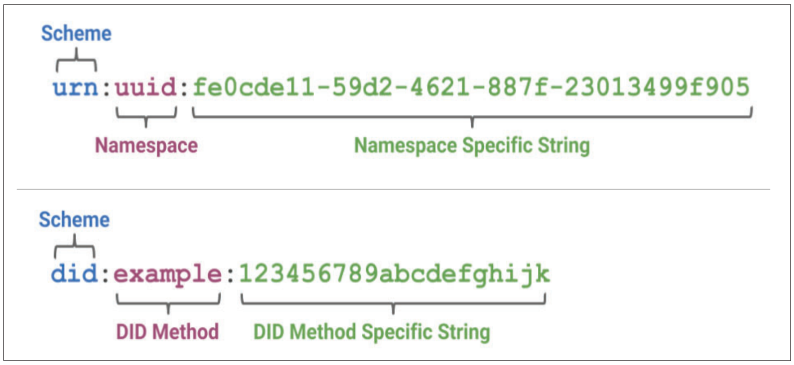
Scheme
urn:uuid:fe0cde11-59d2-4621-887f-23013499f905
Namespace
Namespace Specific String
Scheme
did:example:123456789abcdefghijk
DID Method
DID Method Specific String

FIGURE 2. The structure of DIDs follows the same basic pattern as URNs.

Figure 2 shows how DID syntax resembles URN syntax as defined in RFC 8141.

The DID specification also resembles the URN specification in that it defines a generic URI scheme for defining other specific URI schemes. In the case of DIDs, these are called *DID methods*. Each DID method is defined by its own DID method specification, which must include:

- The target DID network (blockchain, distributed ledger, distributed file system, or other decentralized system) on which the DID

The growing variety of these methods reflects the increasing strength of this foundational layer of the ToIP stack. DIDs anchored at this layer are fundamental to secure DID communications at Layer Two, digital credential issuance and verification at Layer Three, and cryptographically-verifiable governance frameworks at Layer Four.

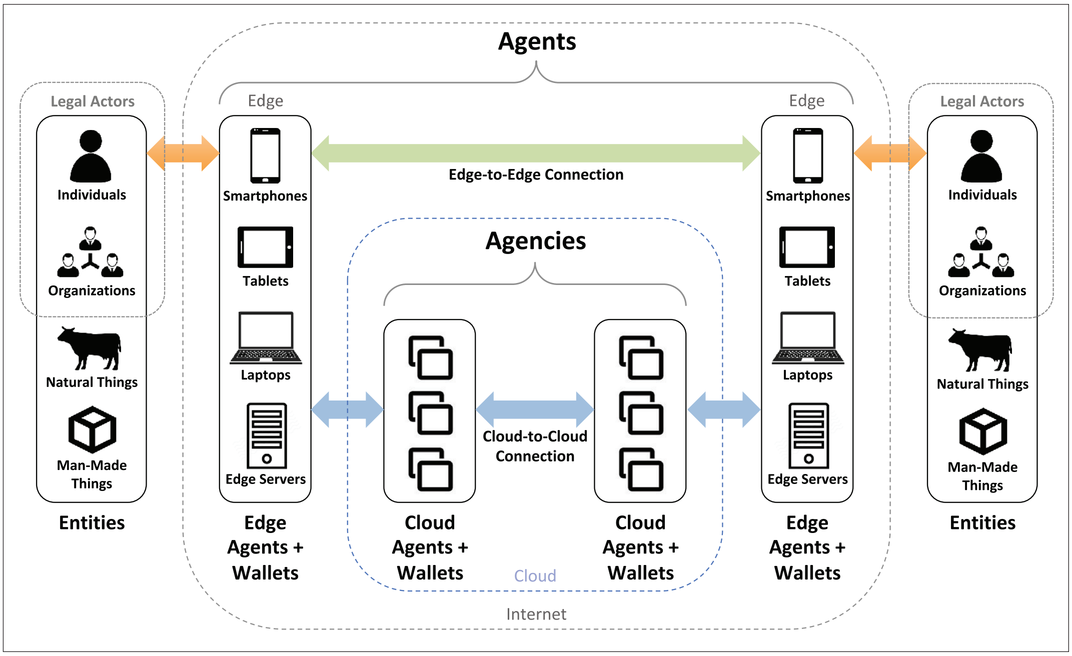

Figure 3. At layer two, agents and wallets communicate using DIDComm standards.

method operates

- The DID method name
- The syntax of the DID method-specific string
- The CRUD (Create, Read, Update, Delete) operations for DIDs and DID documents conformant to the specification

This solution for binding a globally unique identifier to the cryptographic keys and other interaction metadata necessary to prove control of that identifier has proved exceedingly popular. Over 40 DID methods have already been registered in the informal DID Method Registry [6] maintained by the W3C Credentials Community Group. They include methods for:

- Permissionless blockchains including Bitcoin (Three methods), Ethereum (five methods), Veres One, IOTA, RChain, Ontology, and so on
- Permissioned ledgers such as Sovrin
- Distributed file systems such as IPFS
- Ledgerless P2P networks such as git, JLINC, and peer DIDs

The growing variety of these methods reflects the increasing strength of this foundational Layer of the ToIP stack. As we describe below, DIDs anchored at this layer are fundamental to secure DID communications at Layer Two, digital credential issuance and verification at Layer Three, and cryptographically verifiable governance frameworks at Layer Four.

Layer Two: DIDComm

Layer Two of the Trust over IP stack describes the DIDComm messaging standards [7] that establish a cryptographically secured means by which any two software agents (peers) can securely communicate either directly edge-to-edge or via intermediate cloud agents (Fig. 3). What is unique about DIDComm is that the peers who are party to the connection are each individually responsible for:

- The generation of their DID
- The key pairs in a DID document required to establish the secure messaging between them
- The subsequent key rotation or revocation of those keys

This system of pairwise pseudonymous DIDs and keys is specified in the *Peer DID Method Specification* [8].

This layer of the stack strongly differentiates it from most previous trust systems, which rely on some aspect of centralization in terms of identifier creation and control, cryptographic key generation, or both. This architecture is possible because there is now a separation of concerns between the means of establishing a secure communications channel (Layers One and Two) and the means of establishing peer trust (Layer Three, below). DIDComm provides a way for entities to establish permanent peer-to-peer connections without the aid of an intermediary.

At Layer Two, every agent is paired with a *digital wallet* — anywhere from a very simple static wallet to a highly sophisticated enterprise-grade key server — that safeguards sensitive data such as key pairs, zero-knowledge proof blinded secrets, verifiable credentials, and other cryptographic material needed to establish and maintain technical trust. This is one of the fastest evolving components of the ToIP stack. A March 2019 industry study of the current and emerging state of digital wallets by co-author Darrell O'Donnell [9] identified an urgent need for industry standardization, particularly for interoperability in the following areas:

- Backup and recovery (technology and process)
- Credential exchange protocols (see Layer Three, below)
- Certification and accreditation regimes for software and hardware components
- Enterprise use cases including hierarchical agent topologies
- Roles of third parties in agent/wallet usage (e.g., can a financial institution provide digital asset protection without undue access to

digital assets?).

Many agents will also be paired with a *digital hub* — a data store controlled exclusively by a DID subject where all the data is encrypted with the private keys in the subject's digital wallet. A DID subject might have a single data hub or a set of distributed data hubs that automatically stay synchronized according to the owner's preferences. Work on standardizing digital hubs is now proceeding in collaboration across the Decentralized Identity Foundation, the Hyperledger Aries Project, and the W3C Credentials Community Group.

Layer Three: Verifiable Credential Exchange

If the purpose of Layers One and Two is to establish *cryptographic trust* between peers, the purpose of Layers Three and Four is to establish *human trust* between peers — trust between real-world individuals and organizations and the things with which they interact (devices, sensors, appliances, vehicles, buildings, cities, etc.).

The definition of Layer Three conforms very closely to Mr. Cameron's original vision for Information Cards — digital credentials that holders can use to prove claims about their identity the same way we do with the credentials in our physical wallets today. The only real differences are:

- The term used now is *verifiable credentials* rather than Information Cards.
- The serialization format is JSON-LD instead of XML.
- The recommended identifiers for issuers and holders of verifiable credentials are DIDs instead of URLs or X.500 distinguished names (thus enabling decentralized PKI, widely considered vital to broad adoption).

The original work to develop the verifiable credentials model was led by Manu Sporny and David Longley at the W3C Credentials Community Group,[1] which still maintains the Verifiable Credentials Primer [10]. In the spring of 2017 this work was contributed to the newly formed Verifiable Claims Working Group (VCWG). In August 2019 the VCWG finished what is now a full W3C Recommendation: the Verifiable Credentials Data Model 1.0 [11].[2]

Figure 4 is the VCWG's diagram of the three core roles in verifiable credential exchange.

From the standpoint of the ToIP stack, exchange of verifiable credentials is performed by agents using an extension of the DIDComm protocol. This is another area of intense activity in the Hyperledger Aries project, where extension protocol specifications are being published as part of the DIDComm suite [7]. The current credential exchange protocol supports two types of credentials: those that do not use zero-knowledge proof (ZKP) cryptography, which are easily correlatable, and ZKP credentials that enable credential holders to selectively disclose claims to verifiers without correlation — a major advancement in Privacy by Design architecture that is supportive of the EU General Data Protection Regulation (GDPR), the California Consumer Privacy Act (CCPA), and similar data protection regulations.

The power of interoperable verifiable credentials is that they enable any issuer to assert any set of claims to any holder who can then prove them to any verifier. This is a fully decentralized system that works just like the credentials we carry in our wallets today — and thus can be adapted to any set of requirements from any trust community. In fact, in many cases the ToIP stack will not require new "trust infrastructure" at all, but will enable credentials that are currently issued, held, and verified *physically* to be issued, held, and verified digitally.

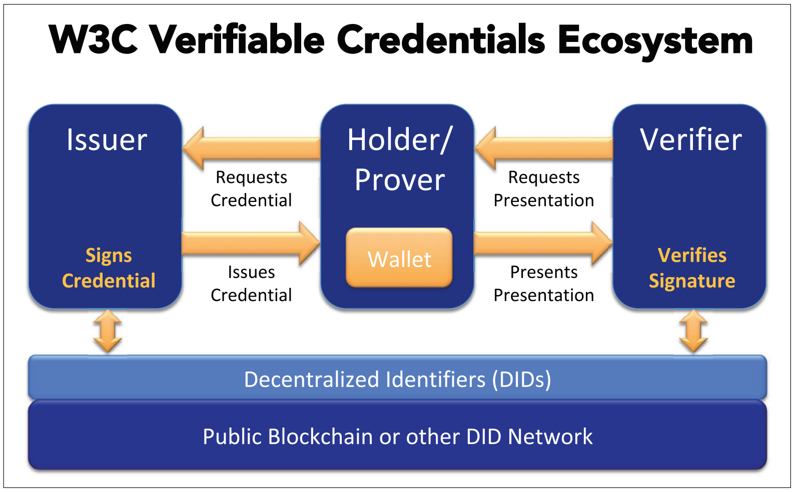

Figure 4. The three core roles in the W3C Verifiable Credentials ecosystem.

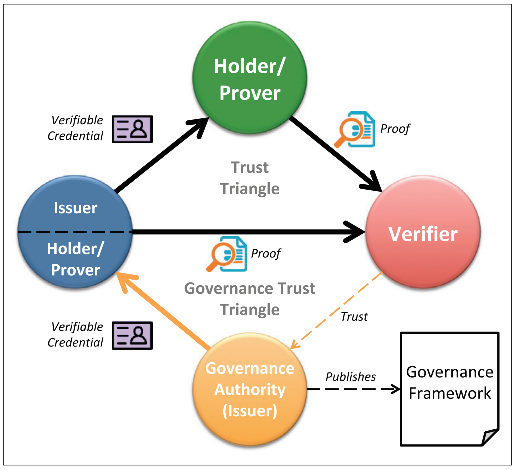

Figure 5. Trust triangles and the role of governance frameworks.

Layer Four: Governance Frameworks

Verifiable digital credentials depend on the same "trust triangle" as physical credentials, as shown in the top half of Fig. 5. But the trust networks and trust frameworks behind the most successful physical credentials — passports, driving licenses, credit cards, health insurance cards — have taken decades to evolve, and in some cases include government regulations.

How can we develop the same degree of trust in digital ToIP infrastructure without having to wait decades for new legislation? The answer is the second layer of trust triangles represented by the bottom half of Fig. 5: *governance frameworks*.

Any group of issuers who wish to standardize, strengthen, and scale the credentials they offer can join together under the auspices of a spon-

[1] Formerly the Verifiable Claims Task Force.

[2] It is a historical artifact of Kim Cameron's and Microsoft's work on claims-based identity that the W3C Working Group is called the Verifiable Claims Working Group while the specification now uses the term verifiable credentials.

Layer	Role	Description
Layer Four: Governance Frameworks	Governance Authority	Specifies a governance framework (GF)
	Auditor	Audits participants for compliance with a GF
	Auditor Accreditor	Accredits auditors for a GF
Layer Three: Credential Exchange	Trust Anchor	Authoritative issuer of a credential under a GF
	Credential Registry	Authoritative holder of credentials for discovery
	Insurer	Insures issuers operating under the terms of a GF
Layer Two: DIDComm	Hardware Developer	Provides ToIP-compliant hardware
	Software Developer	Provides ToIP-compliant edge agents and wallets
	Agency	Hosts ToIP-compliant cloud agents
Layer One: DID Networks	Transaction Author	Initiates a transaction on a DID network
	Transaction Endorser	Facilitates transaction author transactions
	Steward	Operates a node of a permissioned DID network

Table 1. Standard roles in the ToIP governance stack.

soring authority to craft a governance framework. No matter the form of the organization — government, consortium, association, cooperative — the purpose is the same: define the business, legal, and technical rules by which the members agree to operate in order to achieve trust.

This, of course, is exactly how Visa and Mastercard — two of the world's very largest trust networks — have scaled. Any bank or merchant can verify in seconds that another bank or merchant is a member of the network and thus bound by its rules.

With the ToIP stack, this architecture can be applied to any set of roles and/or credentials for any trust community of any size. Table 1 summarizes the 12 standard governance roles defined in the ToIP governance metamodel. Each of these roles is associated with different specific responsibilities in the ecosystem. In most cases a participant will be able to prove what role it is playing in a particular governance framework using a verifiable credential issued by the governance authority or its delegate. Note than an entity may take on more than one role within a given layer based on its goals.

The governance metamodel represented by this half of the ToIP stack is inspired by the Sovrin Governance Framework (SGF) [12]. Two generations of the SGF has been developed over the past three years by the Sovrin Foundation, the governance authority for the Sovrin public ledger for self-sovereign identity (SSI). The SGF currently provides both: a) a Layer One governance framework for Transaction Authors, Transaction Endorsers, and Stewards of the Sovrin ledger, and b) the foundation for a Layer Four identity metasystem. The SGF Working Group is currently working on a third generation of the SGF that will fully incorporate the ToIP stack.

The ToIP governance stack is also designed to be compatible with — and an implementation vehicle for — national governance frameworks such as the Pan-Canadian Trust Framework (PCTF) from the Digital Identity and Authentication Council of Canada (DIACC) [13]. Co-author John Jordan and his team at the Province of British Columbia (BC) have already implemented a verifiable credential registry service called VON (Verified Organization Network) [14] using open source code from the Hyperledger Indy and Hyperledger Aries projects at the Linux Foundation. Between BC and the Province of Ontario, almost 10 million business license credentials have been issued into the VON registries. [15]

Conclusion: A Trust Layer for the Internet

The emergence of the ToIP stack is a watershed moment for the digital landscape: it marks the advent of the solid, decentralized, privacy-respecting trust layer for the Internet that we've been missing for decades. Progress has come by building on the momentum of blockchain technology and developing open standards and open source for decentralized identity and verifiable credentials. Now forward-looking governments and institutions are starting to embrace this solution as a way to protect their constituents, employees, and customers from a growing army of digital predators. ToIP provides an open, neutral, decentralized architecture that can help us transition from the risky "Wild West" Internet of today toward the civilized Internet of tomorrow where trust is the norm and not the exception.

Although building trust at a distance and at scale is a hard problem, so too was the problem of interconnecting disparate local area networks 50 years ago. The TCP/IP stack solved that problem and triggered an explosion of possibility that changed the face of our global economy and society. Now the ToIP stack can solve the problem of trust on the Internet and trigger a second explosion of human possibility.

Future Work

The ToIP stack is still young. Although the architectural distinctions between its four layers are rooted in clean separations based on cryptography, network architecture, and human processes, the specific technical, legal, and business standards at each layer still require further development, standardization, testing, and "hardening" in real-world implementation experience.

In addition, the emergence of the ToIP stack as an identity and credential metasystem does not displace existing identity protocols and systems any more than the emergence of the TCP/IP stack displaced existing local area networks when the Internet was born. Much work is currently ongoing and much more work remains to be done to layer the ToIP stack over and integrate it with legacy identity, authentication, and authorization systems and protocols including SAML, OAuth, OpenID Connect, FIDO, UMA, and other Internet-scale standards.

To facilitate this work, the authors and others are planning to form a new Linux Foundation project whose only mission is coordinating the definition, testing, integration, and adoption of the ToIP stack. This project will collaborate with open source, open standards, and open governance work currently underway at Hyperledger, the Decentralized Identity Foundation, the Sovrin

Foundation, W3C, the Internet Engineering Task Force, and other projects and SDOs working on decentralized identity, security, privacy, and trust infrastructure.

References

[1] K. Cameron, "The Laws of Identity," May 2005; https://blogs.msdn.microsoft.com/tonytri/2008/07/10/the-laws-of-identity/, accessed Nov. 29, 2019.

[2] K. Cameron, "Microsoft's Vision for an Identity Metasystem," Oct. 2005; http://www.identityblog.com/stories/2005/10/06/IdentityMetasystem.pdf, accessed July 6, 2019.

[3] Sovrin Governance Framework Working Group, "Sovrin Glossary Appendix D," Mar. 2019; https://sovrin.org/wp-content/uploads/Sovrin-Glossary-V2.pdf, accessed July 6, 2019.

[4] D. Reed *et al.*, "Decentralized Identifiers (DIDs) v0.13," June 2019; https://w3c-ccg.github.io/did-spec/, accessed July 6, 2019.

[5] W3C Credentials Community Group, "DID Primer," Jan. 2019; https://w3c-ccg.github.io/did-primer/, accessed July 6, 2019.

[6] W3C Credentials Community Group, "DID Method Registry," June 2019; https://w3c-ccg.github.io/did-method-registry/, accessed July 6, 2019.

[7] D. Hardman, "DID Communication," Jan. 2019; https://github.com/hyperledger/aries-rfcs/blob/master/concepts/0005-didcomm/README.md, accessed July 6, 2019.

[8] D. Hardman *et al.*, "Peer DID Method 1.0 Specification," July 2019; https://openssi.github.io/peer-did-method-spec/, accessed July 6, 2019.

[9] D. O'Donnell, "The Current and Future State of Digital Wallets," Mar. 2019; https://continuumloop.s3.amazonaws.com/WalletReport/The-Current-and-Future-State-of-Digital-Wallets-v1.0-FINAL.pdf, accessed July 6, 2019.

[10] M. Sporny, "Verifiable Credentials Primer," Feb. 2019; https://github.com/WebOfTrustInfo/rwot8-barcelona/blob/master/topics-and-advance-readings/verifiable-credentials-primer.md, accessed July 6, 2019.

[11] M. Sporny *et al.*, "Verifiable Credentials Data Model 1.0," Sept. 2019; https://www.w3.org/TR/verifiable-claims-data-model/, accessed Sept. 25, 2019.

[12] Sovrin Governance Framework Working Group, "Sovrin Governance Framework V2," Mar. 2019; https://sovrin.org/governance-framework/, accessed July 6, 2019.

[13] DIACC, "Pan-Canadian Trust Framework," May 2019; https://diacc.ca/pan-canadian-trust-framework/, accessed July 6, 2019.

[14] Governments of British Columbia, Ontario, and Canada, "Verified Organizations Network (VON," June 2019; https://vonx.io/, accessed July 6, 2019.

[15] Sovrin Foundation, "Use Case Spotlight: The Government of British Columbia," Mar. 2019; https://sovrin.org/use-case-spotlight-the-government-of-british-columbia-uses-the-sovrin-network-to-take-strides-towards-a-fully-digital-economy/, accessed July 30, 2019.

Biographies

Matthew Davie (need bio)

Dan Gisolfi (need bio)

Daniel Hardman (need bio)

John Jordan (need bio)

Darrell O'Donnell (need bio)

Drummond Reed (need bio)