# Architecture for Self-Sovereign Digital Identity

Kalman C. Toth
NexGenID
Portland, OR, USA
kalmanctoth@gmail.com

Alan Anderson-Priddy
Portland State University
Portland, OR, USA
andersonpriddy@gmail.com

## Abstract

Our identity architecture significantly improves upon the current patchwork of identity schemes on the web by integrating user authentication with identity specification, virtualization, proofing, and attestation. It mimics how identities are handled in the physical world to provide users digital identities they can use to *prove who they are*. Our *digital identities* are said to be *self-sovereign* because they are tightly controlled within the personal device of the owner. Each digital identity has a *sovereign image* that includes private owner information plus public/private key-pairs used to secure transactions and ensure that owners cannot repudiate their actions when identifying themselves, attesting digital identities and other artifacts, and registering digital identities.

**keywords:** security, digital identity, privacy, authentication.

## 1 Introduction

Our work has been motivated by the challenge of solving the identity crisis [6] [9] caused by excessive centralization of personal information and over-dependency on passwords.

Today, private information is widely scattered to support a patchwork of identity schemes, bridges, add-ons, and protocols for identity access and management. Remote information access and sharing is much too dependent on passwords which can be stolen, lost, cracked, and hacked. Many web-based business models capture enormous volumes of private information while providing inadequate governance, control, and privacy protection. Most consumers/citizens are unaware or oblivious to the risks. The alarming growth of service provider repository breaches, disclosure of private data, fraud, and unattributed information ("*fake news"*), confirm that the identity crisis has not been solved.

Dick Hardt, former member of the OpenID Foundation Board, said in [5] that the Internet needs a "generative" user-centric identity platform whereby users control and use "deep rich digital personas" by means of "a single simple protocol that everyone implements". Expanding upon his vision, we believe digital identities should also be highly intuitive and easy to use, mimicking identity issuance in the real world.

We have examined identity on the web in the context of existing identity technologies and familiar processes used to issue physical identities (e.g. passports, drivers licenses).

Online user access and collaboration continues to be predominately secured by server-centric remote authentication methods including remote access passwords (*what users know*) and biometric authenticators (*what users are*). To overcome the limitations of server-centricity, writers [1] [8] are advocating self-sovereign identity schemes wherein digital identities are strongly controlled by their owners.

In this paper we describe our architecture for self-sovereign digital identity. Referring now to Figure 1, users own virtualized digital identities held within their personal devices which they control by way of password/PIN and/or biometric authentication. Owners specify their identities and can request other parties to proof, attest, and issue them. Both owners and issuers can register digital identities in a "*proof-of-existence*" identity registry. Relying parties can check the veracity of digital identities directly with owners, and/or by way of the identity registry. Registered identities are hashed and stored rendering the identity registry immune to breaches.

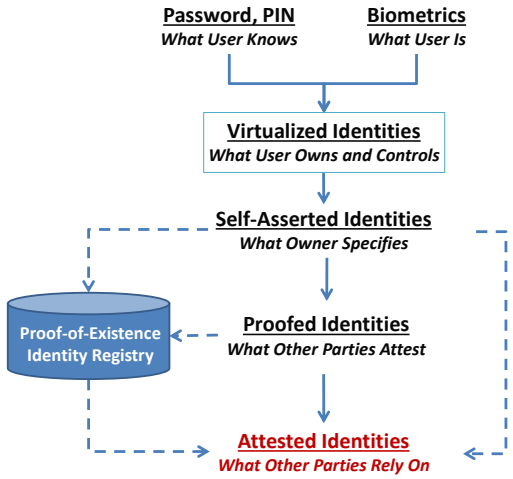


**Figure 1: Core Properties of Self-Sovereign Identity**

## 2 Our Approach to Self-Sovereignty

Our identity architecture for *self-sovereignty* deploys digital identities that are tightly controlled by their owners using their *personal device(s)* (e.g. smart phones, tablets, laptops), each device having a pre-installed *identity engine*. When an owner uses his identity engine to specify a digital identity, a *sovereign image* is created specifying claims, attributes and

Presented at CAINE 2018, New Oreleans LA, October 8-10, 2018, published by ISCA

images characterizing the owner, claims being consistent with Kim Cameron's definition [4]. Instead of entering account names and passwords, users select and present their digital identities to relying parties to identify themselves. Once collaborating parties have reliably exchanged their digital identities they can use them to collaborate securely. Transactions are bilaterally signed and encrypted to thwart phishing, pharming and other impersonation attacks.

As depicted in Figure 2, personal devices collaborate on behalf of their owners to specify and tightly control digital identities to identify owners; verify digital identities; proof, attest and issue identities; register and verify them using an identity registry; notarize documents; reliably transfer identities; and securely collaborate. Our digital identities are said to be "*self-sovereign*" because owners control them throughout their useful lives from when they are created, to when they are deleted or retired.

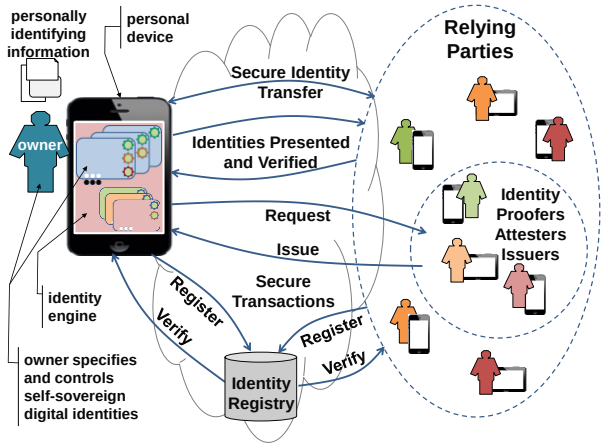


**Figure 2: Collaboration Protocols and Transactions**

Our identity architecture enables device owners to create multiple rich personas, namely self-sovereign digital identities that include attributes and images characterizing the owner.

X.509 digital certificates employed by PKI and PGP use a single key-pair and specify limited information characterizing associated web services, domains, and certificate authorities. Asokan [2] recommends using multiple key-pairs to elevate cryptographic strength and discusses *proof-of-possession*.

Our model enables owners to specify comprehensive digital identities (e.g. attributes, photos, logos) and allocate one or more public/private key-pairs to each identity for signing and encrypting transactions, and for affixing identities and attestations of owners to digital identities and other artifacts.

Self-sovereignty (control) over digital identities is accomplished by designing-in a range of identity assurances including authentication and proofing to achieve persistence, portability, usability, interoperability, and verifiability.

The essential features of our identity architecture are detailed (requisite structures and methods) in US patents.

## 2.1 Identity Model, Persistance and Portability

We employ a common identity data model akin to the approach advocated in [10] for the specification of digital identities comprised of characterizing claims. Our model enables owners to control and access their digital identities persisted in memory, and render them portable for reliable transfer, identification, backup, recovery and escrow.

## 2.2 Usability and Ease of Use

Our design ensures that user interfaces are familiar, unambiguous, and easy-to-use [3]. Users specify virtualized digital identities combining characterizing images (photos, logos) and "*claims*" (attributes) as depicted in Figure 3.

Collaborators can visually inspect and intuitively select their virtualized digital identities and those of others for identification purposes; to attest the identities of others; and to secure transactions while preventing impersonation.

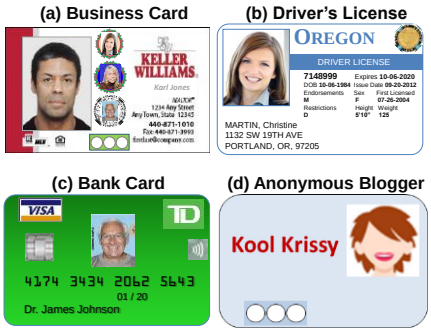


**Figure 3: Virtualized Digital Identities**

Self-sovereign digital identities of owners are potentially operable across a full range of use cases among owners as well as between owners and online web services including: consumer-to-consumer collaboration (email, messaging and conferencing); online access (social, business, enterprise, government, and e-commerce services); and anonymous posting (bulletin boards, blogs, and survey sites).

Figure 4 depicts a user interface where Chris selects Karl's digital business card and one of her own digital identities to launch a Skype collaboration session with Karl.

## 2.3 Controlling Digital Identities

Figure 5 depicts an owner and her personal device with an installed identity engine holding her digital identities as well as those of other parties.

The owner maintains control over the sovereign images of her digital identities by means of the identity engine. Only the owner can use her identity engine to create, store, access, update, expire, delete, and use her digital identities. The owner is authenticated locally. And she can instruct her identity engine to select one of her digital identities, and a

digital identity of another party, to establish a secure session with the identity engine of that other party.


**Figure 4: Skype Session Using Digital Identities**

## 2.4 Locally Authenticating the Owner

Figure 5 also depicts the identity engine controlling the owner's authentication data used by the device's authentication mechanisms to enroll and authenticate the device owner. The identity engine provides a dedicated conduit between the authentication mechanisms and the authentication data; protects this critical data from tampering; and does not reveal this data outside the context of the identity engine. The owner's digital identities, including the owner's authentication data, are thereby strongly bound to the owner and protected from misuse, malware, and surveillance.

The strength of binding depends on the combination of factors used to locally authenticate the owner.

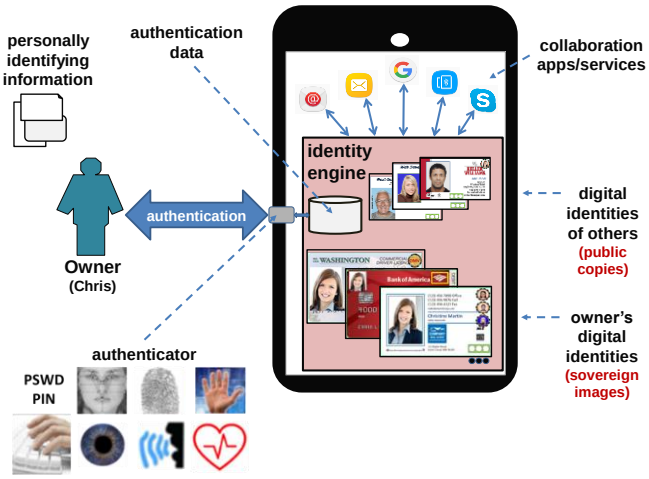


**Figure 5: Owner's Personal Identity Device**

## 2.5 Interoperability Across Identity Layer

Cameron in [4] pointed out that the Internet is crucially missing an identity layer for reliably connecting collaborating parties. Consistent, reliable interoperability can be achieved by establishing a well-behaved identity layer with standard interfaces and services. However, creating such a standard, or standards, will require the application of considerable effort, stakeholder consensus, and time. We will launch a project that progressively deploys personal devices with installed identity engines across an evolving context.

Figure 6 depicts interoperability across the identity layer among device owners, application services, digital wallets, and contact lists. The integrative identity layer between the application layer and the transport layer enables the construction of consistent application programming interfaces supporting collaborative services such as text messaging, email, and conferencing. When launching a collaborative application, the personal device owner uses her identity engine to select digital identities from her wallet and contact list to identify herself and the relying party.

## 2.6 Counterfeit Prevention

Asokan [2] recommends designating multiple public/private key-pairs for distinct purposes to elevate resistance to cryptographic attack. We have adapted his recommendations to thwart the creation of counterfeit (bogus) digital identities.

When the owner specifies a new digital identity, her identity engine creates a master copy of the digital identity called the *sovereign image* which can include multiple public/private key-pairs (see Figure 7). Private (secret) keys are vaulted by the owner's identity engine to protect them from disclosure and tampering by concurrent, and potentially malevolent, software. When the owner selects and presents one of her digital identities to a relying party, the identity engine does not reveal the private keys, delivering only a *public copy of the digital identity*. In other words, the relying party receives only the public keys associated with a presented digital identity.

Depending on the context, risks, number and length of keys, and encryption method(s) used, determining the private encryption key from the paired public encryption key is a hard mathematical problem. Therefore if a malicious party captures the public copy of a digital identity, it is infeasible for that party to discover the private key(s) from the public key(s) to create a counterfeit. Nevertheless, a relying party can use the public keys to challenge an originating owner to determine whether the owner possesses the matching private keys.

## 2.7 Synchronous Verification using Proofs

When establishing a synchronous (interactive) session, collaborating owners play both originating and relying roles.

As illustrated in Figure 7, once an originator has presented the public copy of her digital identity to a relying party, the

identity engine of the relying party can check the veracity of the presented digital identity, and then obtain proof that the originator controls the associated sovereign image. To accomplish this, the relying party's identity engine uses a designated public key of the presented digital identity to execute a *proof-of-possession* challenge [2] which can only be satisfied by using the paired private key of the originator's sovereign image. Such a test determines whether the originating owner's identity engine controls (*possesses*) the matching private key and hence the associated digital identity.

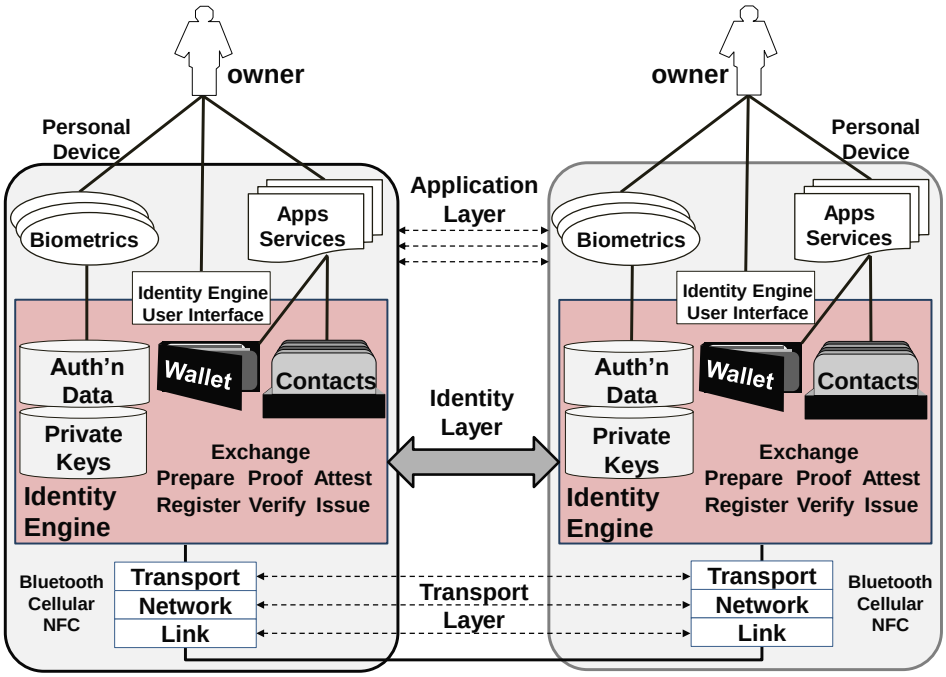


**Figure 6: Application Service Interoperability Enabled by Identity Layer**

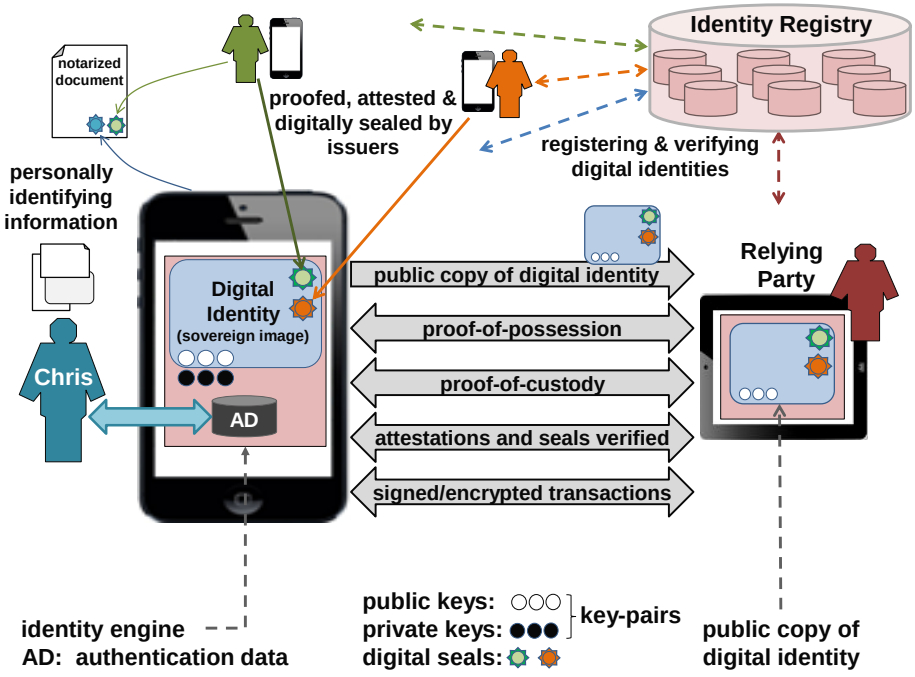


**Figure 7: An Owner, Two Issuers, a Relying Party and an Identity Registry**

If successful, the identity engine of the relying party can send a demand to the identity engine of the originating owner's identity engine to authenticate the holder and send a *proof-of-custody* indication to the relying party's identity engine. Used in combination, these tests determine whether the originator controls the presented digital identity.

Once both parties have identified themselves and have successfully confirmed that the other party has custody of the presented digital identity, neither party can deny having participated in the collaborative session.

## 2.8 Asynchronous Verification using Registry

In support of asynchronous collaboration (e.g. messaging services like email), our architecture incorporates a capability for registering digital identities to enable relying parties to verify acquired digital identities. This mechanism combines our digital sealing method with a *proof-of-existence* method popularized by blockchain [7].

Figure 8 depicts in an identity registry potentially replicated in the form of a distributed database (possibly a distributed ledger using blockchain technology). Each owner's identity engine automates identity registration and identity verification. The figure shows requesting and issuing owners registering digital identities when created and when issued. Relying parties can also use the registry to verify digital identities when presented or acquired.

The registering process hashes the digital identity creating a *hash record*. The registering owner selects one of her digital identities to digitally seal the hash record; link the digital seal to the hash record; and store the hash record, the link, and the digital seal in the identity registry. The digital seal linked to the hash record provides objective evidence that the digital identity was registered by the registering party (owner or issuer). The registering party cannot repudiate having registered the digital identity.

When a relying party has been presented or has acquired a digital identity, he can verify the existence of the digital identity by hashing it and using the hash to locate a matching hash record in the identity registry. If a matching hash record is found, the linked digital seal is verified to determine whether the digital identity was registered by the registering party (the owner or the issuer). The digital identity is valid if it exists in the identity registry and the linked digital seal successfully verifies.

The identity registry can be made publically available because only hashes of registered digital identities are stored, rendering the identity registry immune to breaches.

## 2.9 Mimicking Identity in Physical World

Our architecture mimics identity processes used in the physical world to facilitate user buy-in and adoption. Figure 8(a) illustrates a requester preparing, registering and submitting a digital identity to an issuer who verifies, proofs, attests and issues the digital identity for verification and registration by the requester. Figure 8(b) also illustrates the two parties presenting and verifying each other's digital identities prior to their secure collaboration session.

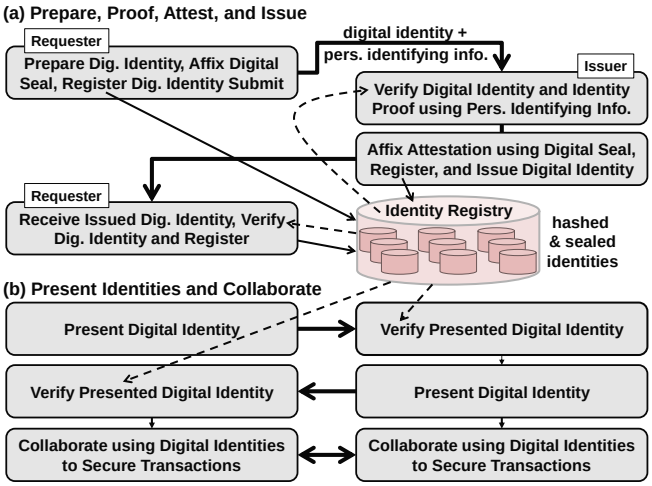


**Figure 8: Registering and Verifying Digital Identities**

## 2.10 Identity Assurance, Proofing, Attestation

Third party identity-proofing and attestation is needed to provide assurances that a digital identity truthfully characterizes the owner and not some other party.

Figure 8(a) illustrates a *requester* registering his digital identity and presenting it to an *issuer* for proofing. Consistent with needs and perceived risks, the requester reveals selected fragments of private and personally identifying information to the issuer.

The identity engine of the issuer then verifies that the requester's digital identity is registered, and conducts identity proofing *in-person* or *online* (when the channel is adequately trusted). If successfully proofed, the issuer selects one of her digital identities to create a digital seal that affixes her identity and an attestation (e.g. "proofed") to the requester's digital identity, registers the attested digital identity, and issues the attested digital identity to the requester. Upon receiving the attested digital identity, the requester acknowledges receipt and registers it.

Identity assurances are elevated for the requester because the issuer cannot repudiate having affixed the attestation to the requester's identity. As illustrated in Figure 8(b), the requester can subsequently present his attested digital identity to relying parties who can verify the affixed digital seal and attestation of the issuer. Relying parties can also use the identity registry to verify the requester's identity.

As depicted in Figure 7, multiple parties can attest, digitally seal, and issue digital identities for an owner. Such *multiple-attestation* incrementally elevates identity assurances associated with an owner's digital identities.

### 2.11 Transferring Digital Identities Reliably

The identity engines of owners can be used to reliably transfer digital identities online by employing the identity registry to verify that they were not corrupted in transit - a reasonable strategy when the risks of a man-in-the-middle attack are low. Collaborating parties can also use their identity engines to securely transfer digital identities by way of in-person encounters using NFC, Bluetooth, WiFi, QR codes, thumb drives, and USB cable.

When owners cannot meet in-person, and online transacting is risky, one of the above techniques can be used to reliably exchange low sensitivity digital identities, subsequently using them to transfer more sensitive digital identities. Another approach is for owners to use their identity engines to exchange a one-time-password or passphrase (OTP) out-of-band (e.g. text, email, or voice), deriving a shared symmetric key which can be used to secure the transfer of the sensitive digital identity.

Mutually trusted passwords over HTTPS, and Diffie-Hellman Key Exchange, can also be employed.

### 2.12 Securing Transactions

Once collaborators have securely exchanged their digital identities, their identity engines can use designated public-private key-pairs of their digital identities to secure their transactions end-to-end, thereby thwarting man-in-the-middle attacks.

### 2.13 Digital Sealing and Notarization

Designated key-pairs of the digital identities of owners can be used to create and verify digital seals used to affix attestations to shared electronic artifacts. When an attestation is affixed to an electronic document (e.g. "this is a true copy"), it has the effect of notarizing the document.

## 3 Concluding Remarks

To facilitate adoption and usability, digital identities are virtualized and handled in a manner that is consistent with identity processes used in the physical world. Encapsulated authentication data (e.g. PINs, biometrics) enable owners to tightly control their secrets and digital identities as well as the acquired public copies of digital identities belonging to other owners. Employing a common data model with integrated public/private keys for structuring digital identities enables persistence and portability. The identity layer supports interoperability with applications and the identity registry, and secures transactions between owners.

Because public/private keys are integral to every digital identity created, owners cannot repudiate having taken critical actions. For example, when presenting, attesting, or registering digital identities, owners cannot deny having done so, thereby providing elevated identity assurances to relying parties. Digital identities can be reliably verified because the private keys of owners' digital identities are not revealed. A relying party can thwart impersonation attempts by obtaining proof-of-possession and proof-of-custody from originators or by verifying digital identities in the identity registry. To create bogus identities the hacker is obliged to successfully break the personal devices and private keys of owners, one device at a time, a prohibitive task.

## 4 Areas for Further Study

We plan to study the following relevant areas: applicability of the Verifiable Claims WG [10] and OpenID [5]; software and protocol correctness; containerization technologies (e.g. Samsung Knox); decentralizing the identity registry; and conducting a vulnerability analysis.

## References


[1] Christopher Allen, "The Path to Self-Sovereign Identity", April 27, 2016, http://www.coindesk.com.

[2] N. Asokan, Baltteri Niemi, Pekka Laitinen, *On the Usefulness of Proof of Possession*, 2nd Annual PKI Workshop, Apr. 28-29, 2013.

[3] Dirk Balfanz, G. Durfee, R. E. Grinter, D.K. Smetters, "In Search of Usable Security: Five Lessons from the Field", IEEE Security and Privacy", Sept./Oct. 2004.

[4] Kim Cameron, *The Laws of Identity*, May 2005, http://myinstantid.com/laws.pdf.

[5] Dick Hardt, *User-Centric Identity, OpenID Foundation Presentation*, Dec. 2010, http://dickhardt.org/2010/12/oidf-2010/index.html#more-85.

[6] Joy Macknight, TheBanker.com, *Will the Digital World Solve the Identity Crisis?*, Jan. 2, 2018, http://www.thebanker.com/Transactions-Technology/Will-the-digital-world-solve-the-identity-crisis.

[7] Kiara Robles, *BlockchainMe, a Tool for Creating Verifiable IDs on the Blockchain*, Dec. 2, 2016, https://github.com/kiarafrobles/blockchainMe.

[8] Sovrin Foundation, "Sovrin: A Protocol and Token for Self-Sovereign Identity and Decentralized Trust", Version 1, January 2018.

[9] Kalman C. Toth, "Brewing Next Generation Identity", Pacific Northwest Software Quality Conference, Portland, OR, Oct. 2015.

[10] World Wide Web Consortium (W3C) Verifiable Claims Working Group, https://www.w3.org/2017/vc/WG.