TechVision
RESEARCH

Decentralized, Blockchain-Enabled Identity Services Gain Traction

Published 05 December 2019

Abstract

On the Internet, “nobody knows you are a dog”. There was and is a fundamental trust problem on the Internet. Enterprises spend trillions of dollars every year trying to properly allocate and remediate business risk based on trusting identity attributes (credentials) presented by employees, customers, and partners. And the problem is ready to take an exponential leap in complexity as billions (some say trillions) of connected devices and services present and consume identity credentials sans human intervention.

Traditional trust models based on the assumption that “the possession and presentation of credentials equals entitlement” are no longer sufficient, and in 2016 we began to see a new Trust over IP (ToIP) model emerge; one based on Decentralized Identity and verifiable credentials. As we began this update to our research on this topic, we wanted to answer these questions:

- Are Decentralized Identity services built on blockchain and leveraging verifiable credentials truly a better model?
- How should our enterprise clients factor these nascent technologies into their strategies and roadmaps?

Based on TechVision’s previous research (Blockchain Identity – 2016, Identity is the New Perimeter – 2017), interviews with key players, the progress we have seen in standards, consortiums and enterprise investments, and most recently face-to-face discussions at our 2019 TechVision Chrysalis conference, we believe Decentralized Identity shows promise as a way provide scalable, peer-to-peer trust across the Internet. On this basis, we recommended that enterprises:

- First get acquainted with the groups working on standards and ecosystems for attribute exchange and interoperability. Also, begin to investigate the early vendor offerings understanding that there will be a lot of volatility.
- Begin mapping you Enterprise IAM architecture and into a capabilities-based framework, such and the TechVision Reference Architecture, and iterate based on requirements, business needs, technology assumptions that will change as decentralized identity is added to the mix.
- “Get you feet wet”; education, pilots, POCs…this is a potential game-changer and enterprises should be well prepared for this new trust model.

Authors:

Gary Rowe
CEO & Principal Consulting Analyst
gary@techvisionresearch.com

Doug Simmons
Principal Consulting Analyst
dsimmons@techvisionresearch.com

Gary Zimmerman
Principal Consulting Analyst
garyz@techvisionresearch.com

© 2019 TechVision Research, all rights reserved

Table of Contents

© 2019 TechVision Research, all rights reserved

Executive Summary

Digital transformation and the proliferation of identities and identifiers are placing new demands on Identity and Access Management (IAM) services. Traditional IAM models are centralized, inflexible and siloed and not easily integrated with many of the newer technologies and business models. These “old school” IAM systems are limiting the ability of enterprises to embrace digital transformation and become a modern digital enterprise.

This report describes a Decentralized Identity (AKA, Self-Sovereign Identity or SSI) movement TechVision sees as an attractive proposition for identity consumers that are overwhelmed by the current approach as well as an opportunity to better mitigate risk for the enterprise. Identity may be the new perimeter, but that doesn’t mean that the enterprise needs to retain so many identifiers and associated personal attributes; they simply need to incorporate distributed identities into their own ecosystems.

The current model for establishing an identity and granting unique credentials is a process that must be replicated for each site an individual connects with. This results in too many IDs and passwords for the average human to remember. To ease this burden, credentials have been shared using federation, Single Sign On (SSO), social login and a myriad of authentication and authorization schemes. But most of these solutions are not solving the fundamental problem—the lack of trust on the Internet.

This lack of trust causes enterprises to spend trillions of dollars every year trying to properly allocate and remediate business risk based on trusting the identity attributes (credentials) presented by employees, customers, and partners. And the problem is ready to take an exponential leap in complexity as billions (some say trillions) of connected devices and services present and consume identity credentials sans human intervention.

Enterprises are looking for a flexible, adaptive and secure IAM foundation and incorporating decentralized identities into this new foundation is a likely outcome—but the challenge is determining when this movement will occur and what the underlying technologies and ecosystems will be.

Decentralized Identity, is a disruptive approach to addressing the trust problem at its core—determining how to prove control, technical trust and human trust as follows:

- The presenter of the attribute (credential) has control of the identifier that the credential was issued to, the credential hasn't been tampered with, and it hasn't been revoked. This establishes technical trust.
- Who issued the credential, what authority do they have to issue it, and what criteria were used to create the credential? This establishes human trust.

While this disruptive approach is still several years away from wide-scale adoption, the impact is so deep and broad that most enterprises should be examining this area and considering how to incorporate a trusted decentralized identity ecosystem within their intermediate to long term planning horizons.

One of the consistent areas of guidance TechVision has given our enterprise clients is that future state IAM services should be flexible, scalable and open/inclusive. This can be achieved with a loosely coupled architecture and incorporating Decentralized Identifiers (DIDs) as a means to allow the owner (technically the controller) of the identifier to gather and maintain attributes – called credentials, associated with their identity. The goal is for specific credentials to be exchanged and verified as required in real-time within the context of the relationship and the transaction being executed. The credential exchange is driven by access policies based on entitlements, the sensitivity of the data, behavior patterns, transactional risk, external verification data and other supporting information.

Identity can be thought of as a new layer in the OSI protocol stack; one that helps to identify users, resources, rights and validate credentials. When we examine the current Internet protocol stack and architecture, we find that virtually everything from physical access to transport, to presentation and applications are represented, but identity and credentials are not explicitly included. TechVision believes that the lack of this "identity layer" is a fundamental flaw that decentralized identity services and/or other identity standards have the potential to remediate.

The Problem; Achieving Trust at Internet Scale

The problem has existed since the early days of the commercialization of the Internet; how to achieve trust at scale in a way that isn't too taxing on the individual and too risky for the organization providing services. We have figured out how to achieve scale via the widely adopted open source implementation of the TCP/IP stack. This offers the capability for any two peer devices to form a connection and exchange data packets regardless of their local network. Without a doubt, the various TCP/IP implementations have driven a tremendous amount of innovation over the last 30 years. However, there remains a significant and widely recognized gap in Internet architecture: a means for peers to establish trust over these digital connections. This gap has often been referred to as "the Internet's missing identity layer". Kim Cameron (former Microsoft Identity leader and global IAM expert) raised this issue a few decades ago and we believe the need for this identity layer is accelerating as individuals and organizations become more and more dependent upon being "always connected" and the complexity of these connections grow.

> *There remains a significant and widely recognized gap in Internet architecture: a means for peers to establish trust over these digital connections.*

This "layer 8 identity service" can provide a means for individuals and organizations to establish and manage the identifiers they own. This capability can be a conduit for a new type of individual control and empowerment that doesn't exist at scale today, and this lack of an identity layer is at least partially responsible for the mess we have today. The challenges aren't related to just a lack of individual empowerment.

Our current IAM ecosystem is enterprise focused meaning that it is viewed as a one-way means for untrusted end-users to prove their "worthiness" in order to access enterprise resources, products and services. This enterprise-focused model is repeated within practically every organization. This fact forces individuals to acquire identities from every enterprise/service provider they interact with, and remember IDs/Passwords, answer security questions, set up MFA across hundreds of sites and applications.

The rub is that the end-user wants a consistent user experience. That does not mean that all end-users have the same user experience, but that a specific end-user wants to use the same identity "agent" over and over for each identity transaction with the enterprises and services he/she interacts with, similar to the interfaces we all see for saving and printing files regardless of the application.

Currently each enterprise (or service provider) provides its own user interface which means the end-user is learning a new interface, sometimes for one-time use (e.g., site registration), sometimes for sporadic use (once a year for electronic tax filing), and on other occasions, every day access (i.e., web-based business applications and mobile apps). This enterprise focus makes the end-user look for shortcuts to create a similar user experience across many enterprises/service providers like:

- Using the same user ID and password across several enterprises/services (for example, an active email address that is easy for the user to remember, and practically guaranteed to be unique).
- Relying on third party Identity Provider's (IDPs) for identity proofing. Again, a single ID and password to provide a consistent experience.
- Using the same "something you know" security questions/answers across multiple enterprises.

And no matter how much the industry preaches password hygiene and rails against user ID and password reuse, **convenience trumps compliance**.

Unfortunately, the bad guys know this all too well. The lack of this consistent identity layer has resulted in individual and business data being spread around like fertilizer on a global field and, as a result, opening up opportunities for compromises at an accelerating level. Breaches and phishing schemes aren't just focused on the targeted enterprise victim, they are focused on using the stolen IDs to leverage access to every other enterprise/service provider that has a relationship with the end-user.

To better visualize the challenge we face on the Internet today, let's think about something we all can relate to; going to a shopping mall. Imagine if we had to prove our identity in order to enter the mall. Then we had to produce (or register) a separate, unique physical ID and password for each store we entered while in that mall. And then we had to produce yet another form of ID (a credit card) in order to complete a transaction within the store. Then imagine the store placing hidden trackers on us (let’s call them cookies, because that doesn’t sound so bad) to watch where we’ve gone and what we looked at, even if we didn’t buy anything. Sounds rather familiar, doesn’t it?

So, what would the above scenario mean for individuals shopping in a conventional mall or shopping center? First, the mall would be a less attractive place to go. Second, it would be hard for every consumer to keep track of the large number physical IDs and passwords that would be required to visit and conduct business in each store. Third, the personal data provided to get the IDs and passwords *from each store* could damage both the individuals that supplied the data and the organizations keeping the data if it were compromised. Ultimately, it sets up an adversarial relationship between the store and the individual as each battles the tension between convenience, security, service, and privacy; all because of a mutual lack of trust.

While the above situation sounds crazy, this is basically what happens in many on-line interactions. It is a fundamentally flawed system at so many levels because individuals that heavily engage online need to remember IDs and passwords from possibly hundreds of sites and share personal information with the owners of each site (i.e., business establishment, social network or service provider) requiring said credentials. Furthermore, these enterprises need to do the same thing with their business partners.

This lack of trust has created a system that:

- Has generated nearly $2 trillion annually in security and compliance remediation costs caused by identity-based fraud.[1]
- Created a yearly $3 trillion data accuracy problem in downstream operations, analytics, and AI applications. [2]
- Exponentially increased attack surfaces through the social and/or proprietary IDP relationships shared and presented across systems today.
- Increased support costs related to two factor authentication, auto-reset procedures, call center support, proprietary hardware fobs that may be required by proprietary IAM solutions.
- Increased compliance costs as regulators weigh in on consumer privacy issues.

Now that we’ve highlighted the problem, the balance of this report will describe a possible viable and scalable solution; Decentralized Identity services and the four-layer architecture of a Decentralized Identity service. We’ll provide a context for our clients to decide if this is the right solution to address the trust problem and when/how an enterprise can move in this direction. We’ll start by describing the basics of Decentralized Identity, then consider a path towards Decentralized

[1] 2018 Cost of Data Breach Study: Impact of Business Continuity Management – Ponemon Institute
[2] Bad Data Costs the U.S. $3 Trillion Per Year – (2018) Harvard Business Review

© 2019 TechVision Research, all rights reserved

Identity, evaluate the state of the industry and conclude with some specific recommendations. We'll now provide a brief tutorial on the basics of Decentralized Identity.

A Decentralized Identity Level-Set

So, how do we solve the challenge of achieving trust at scale across the web? It starts with what Phil Windley described at the TechVision Chrysalis Conference in November 2019 as an "Identity Meta-system". This includes an encapsulating protocol, a consistent user experience and a modular approach that provides user choice. The goal is to have a "life-like" Identity System across the Internet. In looking at identity in this manner we separate the identity component from the rest of the application architecture. This allows us to optimize the end-user experience while minimizing the risks for the enterprise by treating identity proofing as part of the plumbing, not a problem requiring a proprietary solution.

To achieve the above goals, we need to turn the notion of an enterprise centric IAM architecture (I rent you an ID) into a peering paradigm (We share IDs that each of us can prove we control). What does that mean? It means that end-users and enterprises negotiate their digital relationship as peers rather than supplicants or adversaries. That is the underlying foundation of decentralized identity.

> *We need to turn the notion of an enterprise centric IAM architecture (I rent you an ID) into a peering paradigm (We share IDs that each of us can prove we control).*

As Dan Gisolfi, the CTO at IBM described in our 2019 Chrysalis conference panel on Decentralized Identity, initiatives such as the Hyperledger Aries project, the Decentralized Identity Foundation, the Ethereum Foundation, and the W3C Credentials Community Group are all working together to build the standards that define, and the tools to implement the missing identity meta system. Dan described it as "Trust over IP Technology Stack" consisting of two layers focused on technical trust and two layers focused on human trust; all of which are necessary to fully resolve the trust problem. What follows is an excerpt from the currently defined Trust over IP Technology Stack (ToIP) RFC draft being defined by in Hyperledger Aries.

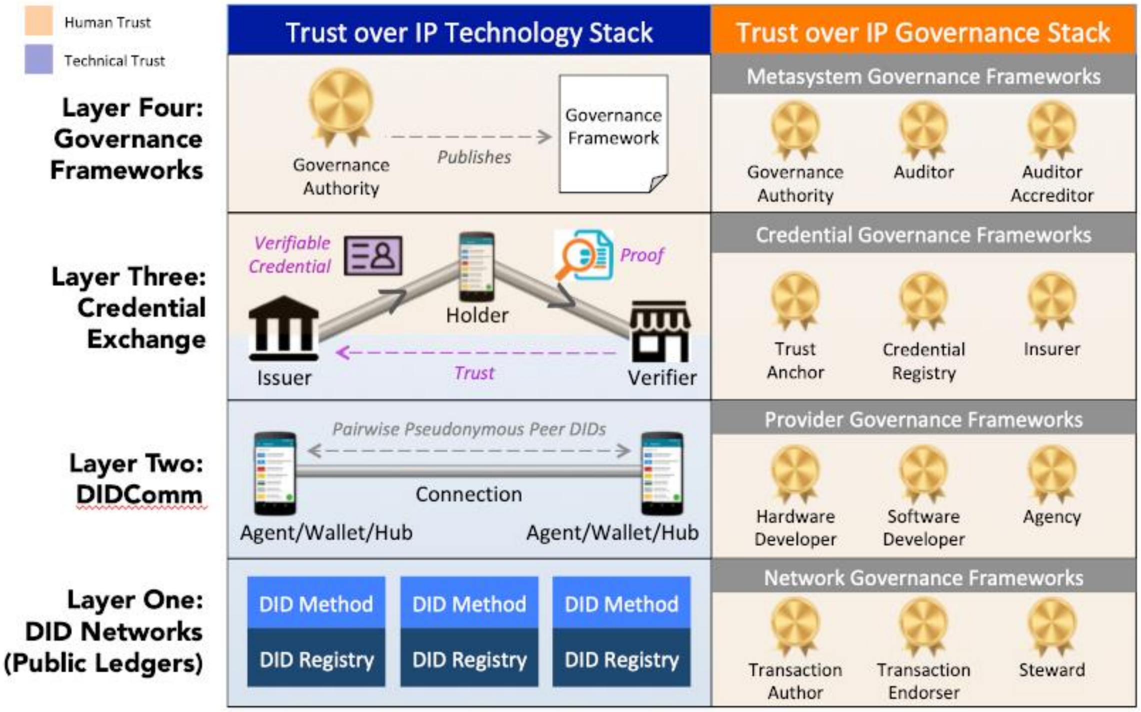

Figure 1 Trust Over IP Stack

It is important to understand that we are using the Trust over IP Stack as an example of a layered architecture and approach to building a Decentralized Identity ecosystem and governance model. There are other approaches and there will be many additional models emerging over the next few years, but these basic layered architectural elements we believe will be a pattern repeated in most models.

Layer One: Decentralized Identifier (DID) Networks

Layer 1 is the base of the Trust over IP stack and is fundamentally made possible by new advancements in cryptography and distributed systems, including blockchains and distributed ledgers. Offering high availability and cryptographic verifiability enables strong roots of trust that are decentralized so there will not be single points of failure. This layer consists of a new type of identifier called a decentralized identifier (DID), a globally unique identifier in which the control of the identifier can be proved using cryptography.

Decentralized Identifiers

Adapting these decentralized systems to be the base layer of the ToIP stack required a new type of globally unique identifier called a Decentralized Identifier (DID). DIDs are designed to provide four core properties:

- **Permanence**. A DID effectively functions as a Uniform Resource Name (URN), i.e., once assigned to an entity (called the DID subject), a DID is a persistent identifier for that entity that should never be reassigned to another entity.
- **Resolvability**. A DID resolves to a DID document—a JSON data structure describing the public key(s) and service endpoint(s) necessary to engage in trusted interactions with the DID subject.
- **Cryptographic verifiability**. The cryptographic material in a DID document enables a DID subject to prove cryptographic control of a DID.
- **Decentralization**. Because they are cryptographically generated and verified, DIDs do not require centralized registration authorities (think central point of failure, compromise, or expense) like those needed for phone numbers, IP addresses, or domain names today.

DID Methods

DID methods are the mechanism by which a DID and its associated DID document are created, read, updated, and deactivated on a specific DID network. Each DID method is defined by its own specification that must include:

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

The target system (called a DID registry) is the environment on which the DID method operates. Note that this is not limited to blockchains or distributed ledgers. DID methods can be designed for distributed databases, file systems, or other system that can serve as a cryptographic root of trust.

DIDs have already proved to be a popular solution to decentralized PKI (public key infrastructure). Over 33 DID methods have already been registered in the informal DID Method Registry hosted by the W3C Credentials Community Group. They include methods for:

- Permissionless blockchains such as Bitcoin (three methods), Ethereum (six methods), Veres One, IOTA, RChain, Ontology, etc.
- Permissioned ledgers such as the Sovrin ledger.
- Distributed file systems such as IPFS.
- Ledgerless Peer-to-peer networks such as git, JLINC, and peer DIDs.

Layer Two: DIDComm

The second layer of the Trust over IP stack is defined by the DIDComm secure messaging standards. This family of Aries specifications establish a cryptographic means by which any two software agents (peers) can securely communicate either directly edge-to-edge or via intermediate cloud agents.

Peer DIDs and DID-to-DID Connections

A fundamental feature of DIDComm is that by default all DID-to-DID connections are established and secured using pairwise pseudonymous peer DIDs as defined in the Peer DID Method Specification. These DIDs are based on key pairs generated and stored by the local cryptographic key management system (KMS, aka "wallet") maintained by each agent[3]. Agents then use the DID Exchange protocol to exchange peer DIDs and DID documents in order to establish and maintain secure private connections between each other—including key rotation or revocation as needed during the lifetime of a trusted relationship. Consider that agents are apps on user devices that act as the principal user interfaces for authentication and authorization to a broad range of unrelated digital sites.

Because all of the components of peer DIDs and DID-to-DID connections are created, stored, and managed at Layer Two, there is no need for them to be registered in a Layer One public DID network. In fact there are good privacy and security reasons not to – these components can stay entirely private to the peers. This also means that, once formed, DID-to-DID connections can be used for any type of secure communications between the peers. Furthermore, these connections are capable of lasting literally forever. There are no intermediary service providers of any kind involved. The only reason a DID-to-DID connection needs to end is that one or both of the peers no longer wants it. We'll next look at how these connections can be accessed and managed via agents and wallets.

Agents and Wallets

At Layer Two, every agent is paired with a digital wallet—or more accurately a KMS (key management service). This KMS can be anything from a very simple static file on an embedded device to a highly sophisticated enterprise-grade key management server. Regardless of the complexity, the job of the KMS is to safeguard sensitive data: key pairs, zero-knowledge proof blinded secrets, verifiable credentials, and the other cryptographic material needed to establish and maintain technical trust. These agents and wallets can also be paired with a hub.

Hubs

Agents may also be paired with a digital hub (something like DropBox, iCloud, or Google drive) – a data store with three special properties:

- It is controlled exclusively by the DID subject (person, organization, or thing) and not by any intermediary or third party.
- All the data is encrypted with private keys in the subject's KMS.

[3] An agent is an addressable network service that serves as a persistent P2P messaging endpoint, coordinates messages and state across multiple clients/edge devices (smartphones, laptops, cars, etc.), maintains an encrypted Key Management backup to simplify key recovery, and simplifies and automates the process of encrypting, storing and sharing data.

- If a DID subject has more than one hub, they can be automatically synchronized according to the owner's preferences.

Work on standardizing digital hubs is proceeding. Next we'll look at a key to supporting trust at scale via Decentralized PKI (DPKI).

DPKI

Encryption, hashing, and digital signatures form the basis of technical trust and these capabilities rely on digital certificates and the established standards for public key infrastructure (PKI). However, the ToIP design eliminates dependence on centralized registries for identifiers as well as centralized certificate authorities for key management – the standard pattern in hierarchical PKI. In cases where the DID registry is a Distributed Ledger each entity may serve as its own root authority – an architecture referred to as Decentralized PKI.

In decentralized PKI, blockchain acts as a decentralized key-value storage. It is capable of securing the data being exchanged to prevent Man-in-the-Middle (MITM) attacks, and to minimize the power of third parties. By bringing the power of blockchain technology to the systems, DPKI resolves the security and control issues associated with traditional PKI systems.

The decentralized nature of the management framework can tackle the problems with the CA systems through certificate revocation, eliminating single points of failure, and reacting fast to misuses of CAs. Blockchain is able to make the process transparent, immutable, and prevent attackers from breaking in, thus effectively avoiding the MITM attacks.

DPKI ensures no third party can compromise the integrity and security of the system as a whole. In blockchain-powered DPKI, the new third parties become miners or validators. The trust is established and maintained based on consensus protocols. Third parties, miners or validators, will have to follow the rules of the protocol, which would limit their roles, and financially reward or punish these third parties to effectively prevent misbehavior in the blockchain.

Layer Three: Verifiable Credential Exchange

Layer One and Layer Two together enable the establishment of cryptographic trust (also called technical trust) between peers. By contrast, 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 (e.g., devices, sensors, appliances, vehicles, buildings, cities, etc.).

The Verifiable Credentials Data Model

Layer Three is currently the most advanced in terms of open standards. After several years of incubation within the W3C Credentials Community Group, the W3C Verifiable Claims Working Group (VCWG) was formed in 2017 and produced the Verifiable Credentials Data Model 1.0, which is currently a W3C Proposed Recommendation. Figure 2 is a diagram of the three core roles in verifiable credential exchange— often called the "trust triangle" as defined by the VCWG.

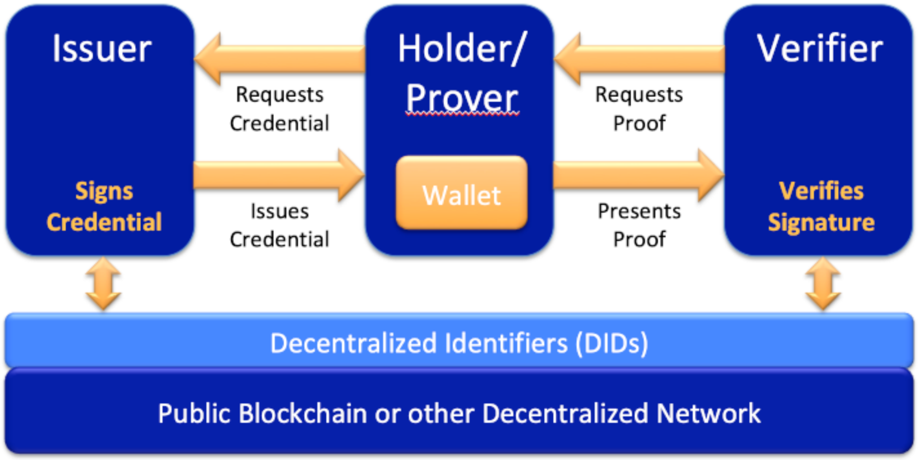

Figure 2 - the Verifiable Credentials Model

With fully interoperable verifiable credentials, any issuer may issue any set of credentials (or claims) to any holder who can then prove them to any verifier. This is a fully decentralized system that uses the same trust triangle as the physical credentials we carry in our physical wallets today. This simple, universal trust model can be adapted to any set of requirements from any trust community. In most cases will not require new "trust infrastructure" at all, but will simply enable existing physical credentials to be transformed into a much more flexible and useful digital format.

Layer Four: Governance Frameworks

The top half of Figure 3 below shows the basic trust triangle architecture used by verifiable credentials. The bottom half shows a second trust triangle—the governance trust triangle – that can solve a number of problems related to the real-world adoption and scalability of verifiable credentials.

© 2019 TechVision Research, all rights reserved

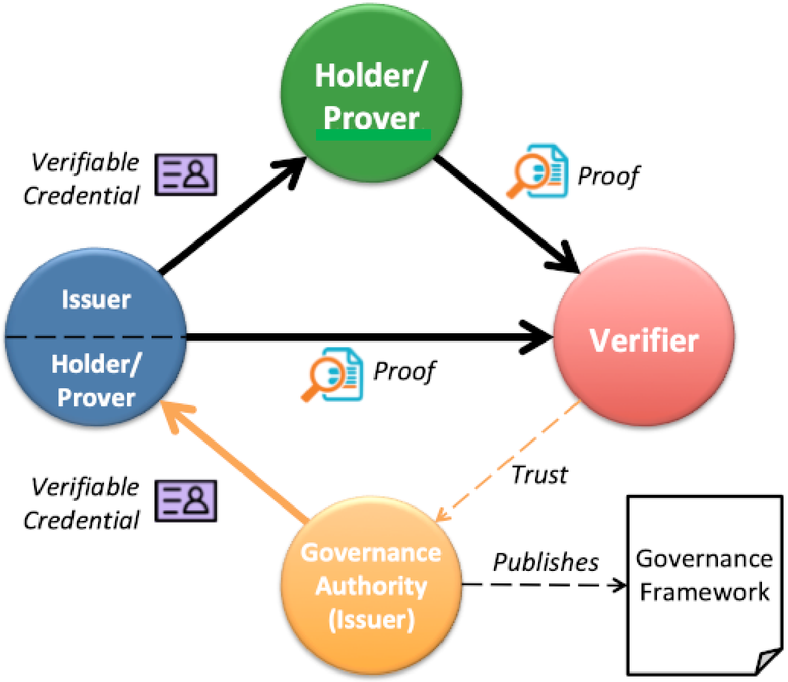

Figure 3 Trust triangles

Governance Authorities

What the governance trust triangle represents is the same governance model that exists for many of the most successful physical credentials we use every day: passports, driving licenses, credit cards, health insurance cards, etc. These credentials are "backed" by rules and policies that in many cases have taken decades to evolve. They are developed, published, and enforced by many different types of governance authorities – private companies, industry consortia, financial networks, and of course governments. The same model can be applied to verifiable credentials simply by having those same governance authorities – or new ones formed explicitly to govern verifiable credentials – publish digital governance frameworks.

The basic principle of ToIP is that end users prove they control their identifier(s) and can present attributes that can be independently verified by the relying party and are provably linked to the identifier they control. Decentralized Identity is an extension of the concept of user centric identity that has been churning around over the past several years.

> *The basic principle of ToIP is that end users prove they control their identifier(s) and can present attributes that can be independently verified by the relying party and are provably linked to the identifier they control.*

About TechVision

World-class research requires world-class consulting analysts and our team is just that. Gaining value from research also means having access to research. All TechVision Research licenses are enterprise licenses; this means everyone that needs access to content can have access to content. We know major technology initiatives involve many different skillsets across an organization and limiting content to a few can compromise the effectiveness of the team and the success of the initiative. Our research leverages our team's in-depth knowledge as well as their real-world consulting experience. We combine great analyst skills with real world client experiences to provide a deep and balanced perspective.

TechVision Consulting builds off our research with specific projects to help organizations better understand, architect, select, build, and deploy infrastructure technologies. Our well-rounded experience and strong analytical skills help us separate the "hype" from the reality. This provides organizations with a deeper understanding of the full scope of vendor capabilities, product life cycles, and a basis for making more informed decisions. We also support vendors in areas such as product and strategy reviews and assessments, requirement analysis, target market assessment, technology trend analysis, go-to-market plan assessment, and gap analysis.

TechVision Updates will provide regular updates on the latest developments with respect to the issues addressed in this report.

About the Authors

Gary Rowe is a seasoned technology analyst, consultant, advisor, executive and entrepreneur. Mr. Rowe helped architect, build and sell two companies and has been on the forefront the standardization and business application of core infrastructure technologies over the past 35 years. Core areas of focus include identity and access management (IAM), Decentralized Identity, Blockchain, Internet of Things, cloud computing, security/risk management, privacy, innovation, AI, new IT/business models and organizational strategies.

He was President of Burton Group from 1999 to 2010, the leading technology infrastructure research and consulting firm. Mr. Rowe grew Burton to over $30+ million in revenue on a self-funded basis, sold Burton to Gartner in 2010 and supported the acquisition as Burton President at Gartner.

Doug Simmons brings more than 25 years of experience in IT security, risk management and identity and access management (IAM). He focuses on IT security, risk management and IAM. Doug holds a double major in Computer Science and Business Administration.

While leading consulting at Burton Group for 10 years and security, and identity management consulting at Gartner for 5 years, Doug has performed hundreds of engagements for large enterprise clients in multiple vertical industries including financial services, health care, higher education, federal and state government, manufacturing, aerospace, energy, utilities and critical infrastructure.

Gary Zimmerman is an experienced executive known for helping companies deliver new offers and expand markets. Accomplishments include launching four companies, 20+ products, building high-performance organizations, and generating millions in sales.

His experience at Neustar, Respect Network, and Sovrin allows him to provide a broad perspective on a variety of subjects including self-sovereign identity, blockchain, enterprise data management, and the data brokerage industry. His experience both enterprise and startup product development give him a unique perspective on innovation.

© 2019 TechVision Research, all rights reserved