The Inevitable Rise of Self-Sovereign Identity

A white paper from the Sovrin Foundation

sovrin

identity for all

Andrew Tobin & Drummond Reed
Foreword by Phillip J. Windley, Chair, Sovrin Foundation

29th September 2016

sovrin.org

Foreword

In 2005, I, along with Kaliya Hamlin and Doc Searls, started a workshop to talk about the problem with identity on the Internet. I think we believed we'd get a bunch of smart people together, hold a few meetings, and get to a solution so we could move on to building the fun stuff. Eleven years and twenty-three workshops later, the Internet Identity Workshop is still going strong and still looking for better solutions to Internet identity.

What makes online identity so hard? Why can't online identity work the way it works in the physical world?

The truth is, we've made a lot of progress since 2005 when most online identity systems were simple username/password schemes and the primary identity tool for most companies was a directory. New requirements like mobile, and new services like social media, have given us identity systems that are much more sophisticated, flexible, privacy-protecting, and user-controllable than anything possible back then.

And yet, we're still a long way from an identity system for the Internet that works the same way the Internet does. The Internet was designed to allow any machine to send messages to any other machine without any administrative authority's permission. In fact, it was designed to route around any attempts to keep those messages from getting through.

An Internet-like identity system would allow any person, organisation, or thing to have an identity relationship (something we call a "claim" in the world of identity) with any other. And to do this without the need for authorisation from someone else. Because anyone can use these identities and the resulting relationships without an intervening authority, they're called "self-sovereign."

An Internet-like identity system has been a long time coming, but I'm excited that recent developments in distributed computing have allowed truly self-sovereign identity systems to be realised. The following paper describes the journey and what awaits us now that we've arrived, and everyone can claim their identity online.

Phillip J. Windley, Ph.D.
Chair
Sovrin Foundation

Table of Contents

- Foreword
- Table of Contents
- Executive Summary
- Terminology Note
- The Internet's Missing Identity Layer
- The Impacts of the Missing Identity Layer
- The Evolution of Internet Identity
  - Centralised
  - Federated
  - User-Centric
  - Self-Sovereign
- The Emergence of Self-Sovereign Identity
- Sovrin: Identity for All
- The Role of the Sovrin Foundation
  - Governance Model
  - The Board of Trustees
  - The Technical Governance Board
  - Executive Director and Staff
  - Sovrin Trust Framework
  - Sovrin Open Source Code
- The Economic Model of Sovrin
  - The Reputation Economy
  - Incentives for Stewards
  - Incentives for Readers and Writers
  - New Global Marketplace for Premium Claims
- Sovrin Infrastructure Providers
  - Stewards
  - Trust Anchors
  - Agents & Agencies
  - Issuers & Claims
  - Relying Parties
- The Road Ahead

Executive Summary

The Internet was built without a standard, explicit way of identifying people or organisations. So websites simply began offering their own local accounts with usernames and passwords, and this has been the predominant solution ever since.

But the Internet has expanded hugely, and people use more and more services daily. This silo-based approach, where users must maintain identities for every site they interact with, has become untenable. It is not just a usability disaster for individuals, it also creates a multitude of data honeypots for hackers—the breach of which compromises trust in all Internet services.

To solve this problem we have tried to connect different identity silos together in various federated models. However these have produced inadvertent side effects such as concentrating control around a small number of providers, increasing data leakage through inadvertent sharing, and raising privacy concerns, all while not actually giving the individual real control.

At the same time, there is a growing economic inefficiency when organisations all around the world have to collect, store and protect the same sort of personal data in their own silos. It is reaching a tipping point.

The next evolution of the Internet will be the creation of a common identity layer that allows people, organisations and things to have their own self-sovereign identity—a digital identity they own and control, and which cannot be taken away from them. Self-sovereign identity is the natural evolution of an ecosystem which has moved faster than its supporting capabilities.

This paper looks at that evolution, and how self-sovereign identity can not only trigger a new wave of innovation, but also provide everyone in the world with a way to establish a portable, secure and controllable identity which is intrinsically theirs.

Terminology Note

All terms used in this and other Sovrin Foundation white papers and documentation are defined in the Sovrin Glossary.

The Internet's Missing Identity Layer

Kim Cameron, Chief Architect of Identity for Microsoft, has said, "The Internet was created without an identity layer." What he meant was that the Internet's addressing system is based on identifying physical endpoints (machines) on a network. People are not endpoints on a network. Therefore, the Internet has no way to uniquely identify people.

Because the Internet can't identify people, websites and applications must do that job. Unfortunately most application developers only concern themselves with their own requirements. This is why usernames and passwords are so deeply imbedded in the fabric of the Web. However it is widely acknowledged that, in addition to terrible user experience, people cannot and will not use usernames and passwords in a way that keeps them safe.

This lack of secure, portable, user-controlled identity has some dire consequences. It means that a person's identity and personal data only exists within the context of each specific website or application he or she uses. Stop using the site or application and the person's digital existence is meaningless. And a user's control over their identity and data must be exerted on a site-by-site, app-by-app basis.

Looking beyond websites and to the broader economy, the global cost to millions of organisations to acquire, store, manage and protect huge volumes of user data is increasing in tandem with the liability associated with holding such data. The sheer inefficiency worldwide data duplication is staggering. Just in the UK alone it is estimated that the cost of identity assurance processes exceeds £3.3bn a year. This would equate to $22bn for the USA if extrapolated to the size of the population[1]. And that's just for initial proofing—the cost of storage, protection, breach and regulation is multiples larger.

The burden of regulation is increasing as user trust decreases. The risk of penetration and breach increases daily, stifling innovation. But to date there has been no credible alternative to the status quo.

The Impacts of the Missing Identity Layer

The impacts of this missing identity layer are becoming acute as more and more of our daily lives depend on digital services.

- Businesses have to develop and manage different security architectures for each platform they deploy (brick & mortar, web, mobile apps).
- CTRL-Shift estimates the total costs of identity assurance processes in the UK exceed £3.3bn. They estimate that this could fall to as little as £150m if people are given control of their own identity data.
- The average retailer cost for each stolen record containing sensitive and confidential information is $165[2].
- 30-40% of contact center call volume is related to password and account recovery[3].
- 25 people in the US fall victim to identity theft every minute—leading to $15 billion in losses from 13.1 million consumers in 2015[4].
- 18% of shoppers abandon their shopping cart due to username and password issues [5].

[1] CTRL-Shift research: Economics of Identity
[2] 2015 Ponemon Institute: Global Cost of Data Breach.
[3] 2011 HDI Support Center Practices and Salary Report.
[4] 2016 Javelin Strategy & Research: 2016 Identity Fraud Study
[5] 2012 Statistica report: Why Do Shoppers Drop Out of an Online Purchase

- 82% of businesses struggle with fake users[6] and on average 10% of a web-facing organisation's user base will be fake.

Without an identity layer, the answer has always been to build an internal database. A silo into which to pour data about customers, whether they like it or not, and whether that data is relevant or not.

Some organisations have "better" silos than others, so convoluted and expensive mechanisms have evolved to pass data from one silo to another, with or without the user's consent, and usually accompanied by unintentional or unwanted data leakage.

The user is the unfortunate victim of this silo mentality. As each organisation spends scarce resources attempting to create the perfect user signup process, all the user sees is yet another web site or app that demands the same details that they entered in the last 50 services they wanted to use. And then they have another username and password to remember.

The user doesn't have their own consolidated digital identity, they just have tens or hundreds of fragments of themselves scattered across different organisations, with no ability to control, update or secure them effectively.

Seen from the user's perspective, the situation is reaching unsustainable levels of absurdity. The constant challenges with usernames and passwords have become cocktail party conversation topics. Seen from the fraudster's perspective, the situation has reached an irresistable level of opportunity.

Into this whirlwind of inefficiency step the regulators. Reacting to the tendency to shuffle data from one silo to another, the use of adhesion contracts[7] against users' best interests, and a frequent lack of explicit informed consent, ever stricter regulation is being imposed. The 261 page EU General Data Protection Regulation[8] is the latest. The intention is to rein in the excesses of an industry that is intoxicated by the promise of ever more data. The immediate consequence will be higher costs for data controllers and data processors as they struggle to comply, adding to the already substantial cost and inefficiency in the identity ecosystem.

Users will bear the brunt of this too, with strange, awkward user experiences as signup processes are patched in the race for compliance. The hugely annoying EU cookie acceptance regulation[9] is nothing compared to what is coming soon.

Something has to change.

[6] 2016 TeleSign report on the Impact of Fake Users.
[7] Adhesion contract: A type of legally binding agreement between two parties to do a certain thing, in which one has all the bargaining power and uses it to write the contract primarily to his or her advantage - West's Encyclopedia of American Law, from The Intention Economy my Doc Searls.
[8] European Parliament General Data Protection Regulation 5419/16 on the protection of natural persons with regard to the processing of personal data and on the free movement of such data..
[9] EU Cookie regulation: Directive 2002/58/EC of the European Parliament. Article 5.3.

The Evolution of Internet Identity

The evolution of internet identity is the result of trying to satisfy three basic requirements:

1. Security - the identity information must be protected from unintentional disclosure;
2. Control - the identity owner must be in control of who can see and access their data and for what purposes;
3. Portability - the user must be able to use their identity data wherever they want and not be tied into a single provider.

In his excellent article "The Path to Self-Sovereign Identity"[10], Christopher Allen provides a clear dissection of the online identity landscape and charts its evolutionary path. His analysis describes 4 stages of development.

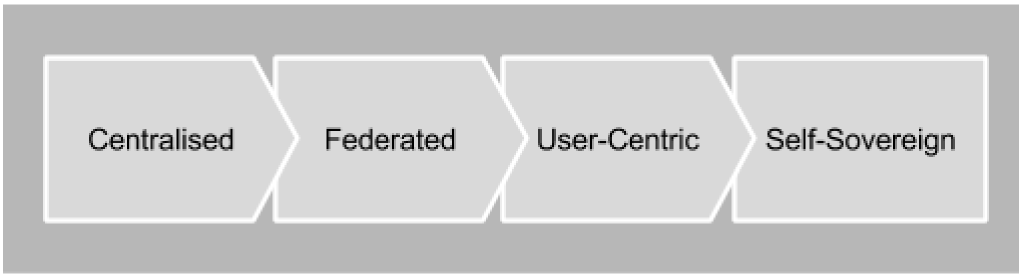

Fig 1. The evolution of online identity

Centralised

The vast majority of Internet identities are centralised. This means that they are owned and controlled by a single entity, such as an eCommerce website or a social network. Within its own domain, centralised identity works fine, but it has struggled to keep pace with the rapid growth and variety of online websites and services with which today's users interact.

Most seriously, because the user doesn't own their identity record, it can be taken away at any time. Language such as the following pervades major website terms and conditions:

> *Yahoo may, without telling you, immediately cancel or limit your access to your Yahoo accounts, certain Yahoo Services and any associated email addresses...*[11]
>
> *If you violate the letter or spirit of this Statement, or otherwise create risk or possible legal exposure for us, we can stop providing all or part of Facebook to you.*[12]
>
> *We reserve the right to modify or terminate the Service or your access to the Service for any reason, without notice, at any time, and without liability to you*[13] *[Instagram]*

[10] The Path to Self Sovereign Identity - Christopher Allen 2016
[11] Yahoo terms of service 2016
[12] Facebook terms 2016
[13] Instagram terms of use 2016

> *We reserve the right at all times (but will not have an obligation) to remove or refuse to distribute any Content on the Services, to suspend or terminate users, and to reclaim usernames without liability to you.*[14] *[Twitter]*

Because the only identities most people have online are centralised, the removal or deletion of an account effectively erases a person's online identity which they may have spent years cultivating and may be of significant value to them, and impossible to replace.

Federated

Federation has been one answer to some of the problems of centralisation. At its simplest, federation gives a degree of portability to a centralised identity, for example enabling a user to login into one service using the credentials of another. At a more complex level, it can allow different services to share details about the user.

Federation is common within large businesses, where single sign-on mechanisms allow a user to access multiple separate internal services such as HR, payroll etc, with a single username and password. In the consumer Internet, federation is visible in services such as Facebook Login, where websites enable users to create accounts and sign in using their Facebook credentials.

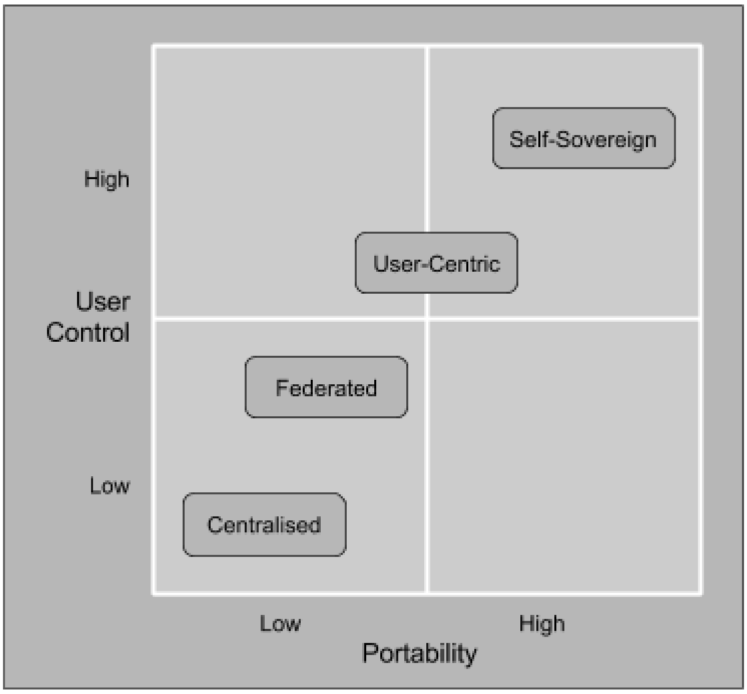

Fig 2: The four stages of online identity (from Christopher Allen #1) against at the axes of portability and control

High assurance federation is emerging as a new market in its own right. The UK government has outsourced digital identity proofing to a number of identity providers who check a user's details and, if they pass the relevant tests, create an account for them. The user can then use this account to log into a number of government services that support the scheme, called "GOV.UK Verify".

Although federation provides a semblance of portability, the power still remains with the identity provider who sits at the centre of the federation web. In fact the implications to a user of having their centrally federated account deleted or compromised are much more profound if that account is their key to many other 3rd party services.

[14] Twitter terms of service 2016

User-Centric

In a prescient paper from 2008, Kim Cameron, Reinhard Posch and Kai Rannenberg describes “A User-Centric Identity Metasystem”[15]. It details an abstracted design for a system which puts the user in control of their own data, the accumulation of that data, and its release to third parties.

In this paper, the authors state that “The core requirement for user control is that the flow of information from Claims Providers to Relying Parties only happens at the request of the user.” The individual fills their own data store with information which they are then able to give their permission to provide to other organisations, keeping a record as they do so.

User-centric identity is most frequently manifested in the form of independent personal data stores at one end of the spectrum, and large social networks at the other end. However the entire spectrum still relies on the user selecting an individual identity provider and agreeing to their often one-sided adhesion contracts[16].

Some existing user-centric implementations are also susceptible to charges of unintended data leakage as they move data from one silo to another, trading the user’s willingness for increased convenience with the exchange of their personal data to a 3rd party. Because they are profit-driven businesses, the user becomes a product to be bought and sold, compromising independence and restricting true portability.

Independent personal data stores also exist, moving more closely towards the vision articulated by Cameron et al in their 2008 paper. They rightly promote the values of individual control, permission and consent, and provide very effective user interfaces. The problem remains that, in a mature personal data store ecosystem, relying parties will need to connect to many such providers to reach a wide customer base, resulting in complex and time consuming integration without economies of scale.

Self-Sovereign

Self-sovereign identity is the final step in this evolution. It provides all three required elements: individual control, security, and full portability. It removes centralised external control aspects from the three previous phases above. The individual (or organisation) to whom the identity pertains completely owns, controls and manages their identity. In this sense the individual is their own identity provider—there is no external party who can claim to “provide” the identity for them because it is intrinsically theirs. The individual’s digital existence is independent of any single organisation. Nobody can take your self-sovereign identity away from you.

The best way to think of self-sovereign identity is as a digital record or container of identity transactions that you control. You can add more data to it yourself, or ask others to do so.

[15] A User-Centric Identity Metasystem October 2008: Cameron, Posch, Rannenberg
[16] Adhesion contracts definition.

You can reveal some or all of it some of the time or all of the time. You can record your consent to share data with others, and easily facilitate that sharing. It is persistent and not reliant on any single third party. Claims made about you in identity transactions can be self-asserted, or asserted by a 3rd party whose authenticity can be independently verified by a relying party.

In a 2016 article[17], Phil Windley describes self-sovereign identity as an "Internet for identity" which, like the Internet itself, has three virtues: no one owns it, everyone can use it, anyone can improve it. Also like the internet, the move to self-sovereign identity is a move from a silo mentality to a layer mentality (Fig 3).

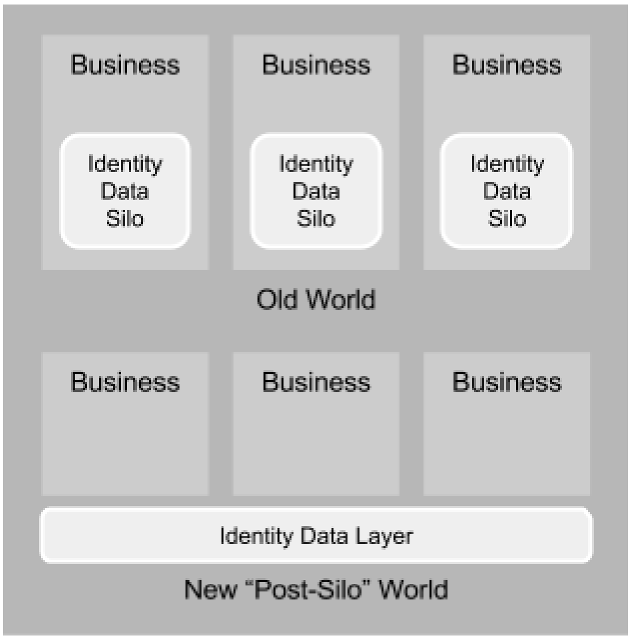

Fig 3: Identity Layer v Identity Silo

Conventional approaches to identity have always struggled with portability, precisely because data is held in closed silos. Complex mechanisms and competing standards have emerged to move data from one silo to another, but none have gained significant traction. While silos persist, individuals will remain dependent on the organisation that own and manage their data for them.

Flourishing use of self-sovereign identity ushers in the post-silo world. Rather than every organisation maintaining their own siloed store of user data and possibly a suite of APIs to connect to other such silos, each organisation can have one single connection to the Internet's identity layer, and immediately benefit from all the organisations that are already present. If a thousand or a million new organisations join the network, no additional work is required to benefit from their presence.

The Emergence of Self-Sovereign Identity

To create the long-missing identity layer of the Internet, a new, trusted infrastructure is required which enables identity owners to share not only identity, but also verified attributes about people, organisations and things, with full permission and consent.

For identities to be truly self-sovereign, this infrastructure needs to reside in an environment of diffuse trust, not belonging to or controlled by any single organisation or even a small group of organisations. Nobody can "turn the lights out". Distributed ledger technology (DLT) is the breakthrough that makes this possible. It enables multiple institutions, organisations and governments to work together for the first time by forming a decentralised network much like the Internet itself, where data is replicated in multiple locations to be resistant to faults and tampering. While distributed ledger technology has been around for some time, new

[17] An Internet for Identity - August 2016: Phil Windley

DLTs such as Bitcoin and Ethereum have resulted in a greater realisation of its potential, particularly with respect to decentralisation and security.

When combined with distributed key management and peer-to-peer sharing of encrypted claims, DLT is what finally makes self-sovereign identity possible. Within this identity layer, mechanisms for discovery, routing of requests, exchanging of data and recording events can exist pervasively, with no single entity being in control.

To define "self-sovereign identity", Christopher Allen's Ten Principles of Self-Sovereign Identity[18] can be summarised in 10 words grouped into three sections:

Security the identity information must be kept secure	**Controllability** the user must be in control of who can see and access their data	**Portability** the user must be able to use their identity data wherever they want and not be tied to a single provider
Protection	Existence	Interoperability
Persistence	Persistence	Transparency
Minimisation	Control	Access
	Consent	

The power of a decentralised platform is well articulated in the 2016 Caribou Digital/Omidyar paper[19] on private sector digital identity in which the authors describe the following:

> *"Open, decentralized systems enable individuals to fully own and manage their own identities, leading to the idea of "self-sovereign" identity systems. These systems use combinations of distributed ledger and encryption technology to create immutable identity records. The individual creates an identity "container" that allows them to accept attributes or credentials from any number of organizations, including the state, in a networked ecosystem that is open to any organization to participate (e.g., to issue credentials).*
>
> *Each organization can decide whether to trust credentials in the container based on which organization verified or attested to them; in other words, a mortgage company may accept a credential issued by a leading global bank, but not one issued by a local bank. Importantly, this model does not require a state-based credential to be initiated (the state credential can be added at a later time, or not at all), which removes a barrier to adoption."*

[18] The Path to Self Sovereign Identity - Christopher Allen 2016
[19] Private-Sector Digital Identity in Emerging Markets - Caribou Digital/Omidyar 2016

In their seminal 2016 paper on the state of the digital identity marketplace[20], Consult Hyperion refer to a model they term “No IDP”, where there is no centralised identity provider (“IDP”). Instead there is a “citizen digital identity controlled by [the] individual”. Acknowledging the main risk that the “technology promises a high degree of privacy but is unproven”, they go on to say:

> *“What Blockchain does is raise the bar for what is viewed as privacy enhancing. In particular, the starting point for digital identity in blockchain is an unlinkable secure identity (a cryptographic public/private key pair) that is only linked to transactions, digital assets or other data with explicit user action. There is also no restriction on the number of identities an individual may choose to create, potentially allowing the user to prevent unwanted linking of transactions.”*

In such a model, the organisations who currently serve as “identity providers” still have a role—it just evolves to the role of “identity proofer”. In other words, someone still needs to carry out checks that an individual is who they say they are as they bridge the physical and digital worlds. Standards of assurance still need to be met and relying parties still need to be able to trust such assurances. What will change will be the ease and simplicity of doing such checks in an environment where the majority of individuals can simply give permission to access verified data to satisfy any given criteria, and the fact that individuals are not beholden to any one bastion of trust in order to have a digital identity.

To summarise, if self-sovereign identity is the natural culmination of the evolution of Internet identity, then delivering it requires diffuse trust, portability, security and no single point of failure or control. Cryptographically secure distributed ledger technology provides the mechanism to make that happen. The next section describes how this capability can be realised.

Sovrin: Identity for All

Creating a new public utility—precisely what the “Internet for identity” should be—requires a way to store identifiers, keys, pointers and proofs without relying on centralised authorities. Using such a mechanism, an individual or organisation can build up a sequence of identity transactions which can reliably prove their identity.

DLT has the capability to serve this purpose. Distributed ledgers are typically split into two types—“public” and “private”—each with benefits and drawbacks (discussed in detail in the Sovrin publication “The Technical Foundations of Sovrin”).

In early 2016 it became clear that a hybrid model - a “public permissioned” ledger - could deliver the best of both worlds; public access with trusted governance. This was the genesis of the Sovrin Identity Network (“Sovrin”), a public permissioned ledger for self-sovereign

[20] Digital Identity Issue Analysis - Consult Hyperion 2016

identity, and the Sovrin Foundation, a global non-profit organisation whose sole purpose is the governance of this ledger and its surrounding ecosystem.

Sovrin puts people, not the organizations that traditionally centralize identity, in charge of decisions about their own privacy and disclosure. This, and the combination of the pervasive nature of the network and the open source nature of its code, enables all kinds of rich innovation: link contracts, revocation, novel payment workflows, asset and document management features, creative forms of escrow, curated reputation, integrations with other emerging technologies, and so on.

Sovrin uses open-source distributed ledger technology. These ledgers are a type of cryptographic database that is provided cooperatively by a global pool of participants instead of a single giant database with a central administrator. In Sovrin, identity records live redundantly in many places, and accrue in transactions orchestrated by many machines. It is protected by strong, industry-standard cryptography and best practices in key management and cybersecurity. The result is a reliable, public source of truth under no single entity's control, robust to system failure, resilient to hacking, and highly immune to subversion by hostile entities.

Sovrin utilizes a "public permissioned" distributed ledger design (Fig 4). The easiest analogy is to the global ATM network: anyone can use an ATM (public), but only those who've been given special permission can add a new ATM to the ATM network (permissioned). With the Sovrin Identity Network, it is the Sovrin Foundation that grants permission for "nodes" (akin to ATMs in the metaphor) to join the network.

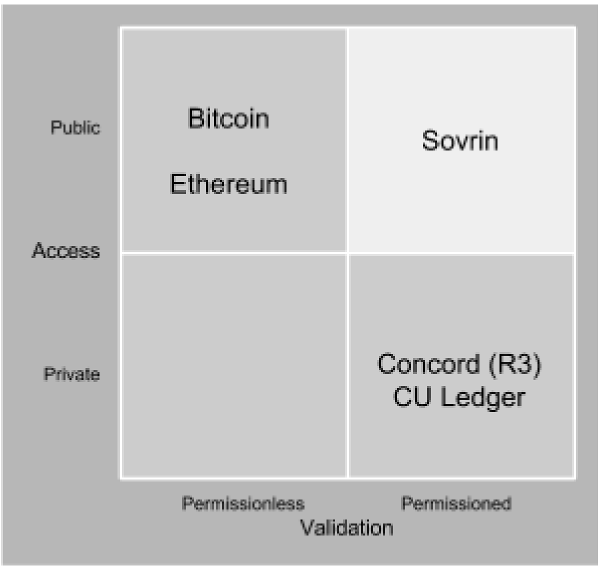

Fig 4. Positioning Sovrin in the Distributed Ledger World

This is a very different approach to "permissionless" ledgers such as those used by Bitcoin and Ethereum, where anyone can join the network. It is also distinguished from private ledgers such as Concord from R3 because Sovrin is intended to be publicly accessible rather than just by a private community of members. Note that "publicly accessible" means Sovrin is open for all to use, it does not mean that all Sovrin identity data is public. In fact the opposite is true: Sovrin identity data is private and can only be shared under the consent of its owner.

A public permissioned ledger is the only option that can achieve both high trust and global adoption. No "mining" is required so it does not need immense computing power and can run at a much higher throughput than permissionless blockchains. It also avoids subversion by any party who can gain a majority of the mining resources (a scenario in which whoever builds the greatest hashing power and has the lowest energy costs controls the ledger). But it is still publicly available to all, unlike private permissioned models such as R3 that can achieve similar efficiencies but only to a restricted population of members.

Another advantage of a public permissioned ledger is that permissioning ensures it can be “fit for purpose”, i.e., dedicated entirely to self-sovereign identity, rather than a generalised “do anything” ledger. This enhances trust. A dedicated, permissioned, purpose-built DLT prevents an organisation or individual from creating a “rogue” application which causes problems affecting all other ledger users. Should anyone attempt to manipulate the Sovrin ledger contrary to the policies laid out by the Sovrin Foundation, the Foundation can agree the relevant sanctions that should be applied through the governance process.

Lastly, a public permissioned ledger designed exclusively for identity has security, privacy, and resiliency properties similar to the Internet itself. The distributed architecture enables high survivability in the event of an attack—and when compared to a centralised data silo, the attack surface of such a ledger is far smaller. Since every transaction can be encrypted individually, the economics of hacking the ledger are also slewed by the time and complexity of mounting an attack on a transaction which may contain no useful information whatsoever. There are no backdoor ways in—administrative access to the data store does not confer access to the data itself—so there is no potential of an “insider” attack from disgruntled employees.

These and other technical benefits are discussed in detail in the Sovrin white paper “The Technical Foundations of Sovrin”.

The Role of the Sovrin Foundation

The creation of a public permissioned ledger means that someone needs to provide the “permission”. Who gives this permission, and how does this governance work?

This highlights the single most important feature of the Sovrin architecture: the human governance layer. The developers of Sovrin believe this layer should be provided by a private-sector international non-profit organisation, representative of all who have Sovrin identities, in much the same way as the Internet’s governance is overseen by ICANN.

For identities to be truly self-sovereign, the nodes of the Sovrin Identity Network cannot be owned or controlled by any single company, organisation, or government. Instead they must be distributed around the world, operated by trusted institutions in multiple countries and industries, forming a decentralised network much like the Internet itself. And like the Internet, they must all run the same distributed consensus protocol that has been publicly vetted for security and privacy.

This is the purpose of the Sovrin Foundation, a not-for-profit global consortium dedicated exclusively to the governance of the Sovrin Identity Network. It’s structure must reflect the same diffuse trust model, i.e., it must represent not only individuals around the world, but the trusted institutions running Sovrin nodes, such as universities, financial organisations and hospitals, as well as issuers of public-sector verifiable attributes such as driving licenses, passports and birth certificates.

Governance Model

The Sovrin Foundation is being launched with a bicameral model governance model representing the two major stakeholders in Sovrin identity infrastructure:

1. **Members** are individuals who have registered identity records on the Sovrin ledger.
2. **Stewards** are trusted institutions who support the goals of the Sovrin Foundation, meet its requirements, and have permission to operate Sovrin ledger nodes.

This governance layer is “as thin as possible but no thinner”, i.e., it establishes only the policies needed for the cooperating nodes of the Sovrin Identity Network to securely store and serve self-sovereign identity records which respect the privacy of each identity owner. All other policies for using Sovrin identity records should be established by communities of participants operating under their own higher level trust frameworks.

The Foundation is charged with carrying out four duties:

1. Develop and maintain the Sovrin Trust Framework—the legal, business, and technical rules governing the selection and monitoring of Sovrin stewards (the legal entities who have permission to operate Sovrin ledger nodes) and operation of the Sovrin ledger.
2. Coordinate and monitor stewards to ensure the ledger is stable, correct, and trustworthy.
3. Govern the Sovrin Project - the open source code that operates, validates, and provides access to the ledger.
4. Promote and advocate the Sovrin Identity Network for self-sovereign identity.

The Board of Trustees

The Sovrin Foundation Board of Trustees (BoT) consists of no less than nine individuals whose duty is to represent the interests of Sovrin identity owners. The founding trustees were selected to represent the balanced interests of global stakeholders. These initial trustees will have a term of one year and will determine the process by which future trustees will be elected.

The BoT approves the overall direction for the Foundation, sets policy, oversees staff, and ensures the financial health of the organisation. The BoT will initially be focused on the steps necessary to move Sovrin from the sandbox stage to a full production network, including development of the Sovrin Trust Framework (below).

The Board also will develop a communications strategy to build global awareness, support and adoption, and create a sustainable funding strategy for the ongoing organisation. It encourages support from individuals and organisations willing to act as trustees, stewards, open source contributors, technical evangelists and global ambassadors.

The Technical Governance Board

The technical expertise required to govern Sovrin is represented by a second board, the Technical Governance Board (TGB). The TGB reports to the Sovrin BoT (Fig 5), and the chair of the TGB is an ex officio member of the BoT. The TGB's responsibilities include the overall technical architecture, strategy and roadmap, threat and risk assessment and planning, guardianship of standards, and management of the Sovrin Project open source code.

The TGB consists of no less than nine individuals whose duty is to represent the interests of Sovrin stewards. The members will be experts in cryptography, security, privacy and distributed computing drawn from three sources:

1. Expert technical staff of the Sovrin Foundation.
2. Technical representatives of Sovrin Stewards.
3. Other 3rd parties with specific technical expertise as deemed appropriate.

The BoT shall invite the initial TGB members and determine the process by which future TGB members will be elected by Sovrin stewards.

Executive Director and Staff

The Executive Director (ED) and staff will be responsible for day-to-day management of the Foundation's activities. They will work with the BoT and TGB to set direction, formulate new policies, and carry out activities that advance the work of the Foundation.

Sovrin Trust Framework

The Sovrin Trust Framework is the core legal and technical governance document for the Sovrin Identity Network. It defines the rights and responsibilities of each of the participants in the network, including Sovrin stewards, trust anchors, and identity owners (see the Sovrin Glossary). Once completed, it will serve as the contract to which all participants must agree.

The Sovrin Foundation BoT has charted the Sovrin Trust Framework Working Group to begin this work. Please contact us if you would like to participate in this effort.

Sovrin Open Source Code

As a public permissioned ledger that serves as a global public utility, the Sovrin Identity Network needs to operate on fully public peer-reviewed open source code. To begin this process, the original source code for Sovrin has been gifted to the Foundation by Evernym Corporation. The Sovrin Project code base consists of three main elements:

1. **Plenum**, the source code which manages how Sovrin validator nodes reach consensus on the ledger.

2. **Sovrin**, the source code which manages the identity and attribute management functions and data structures within the network.
3. **A reference client** which demonstrates how to develop applications for the network.

The Technical Governance Board will be developing guidelines and processes for the Sovrin open source project so it is truly representative of and responsive to the Sovrin community. The Sovrin open source code base is available under the Apache2 license, a permissive open source license that removes obstacles to widespread implementation.

The Economic Model of Sovrin

The Reputation Economy

Reputation has been shown to be a very powerful driver of behaviour. The Sovrin Identity Network enables organisations and individuals to establish, maintain, enhance and communicate their reputations in new and innovative ways.

As an individual's or organisation's Sovrin identity builds up over time, so does their reputation. Stepping up from a low trust level to a higher trust level happens seamlessly as more verified attributes and claims are accumulated by the identity owner. This reputation becomes an asset of the identity owner. For example, an individual may chose to reveal their reputation to others to establish and reinforce trust, or an organisation may publish its Sovrin-based reputation ratings as a badge of honour.

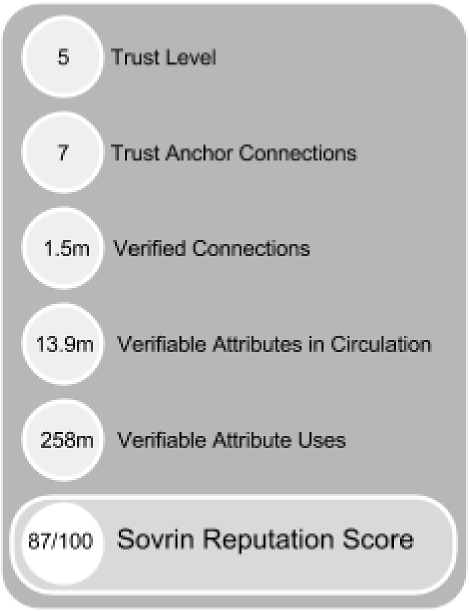

5 Trust Level

7 Trust Anchor Connections

1.5m Verified Connections

13.9m Verifiable Attributes in Circulation

258m Verifiable Attribute Uses

87/100 Sovrin Reputation Score

Fig 6. How an organisation's reputation score could look.

This also produces a virtuous network effect. Organisations that are trusted by other organisations as providers of verified claims automatically enhance their own reputations. The more individuals and organisations that rely on your claims, the higher your reputation.

Sovrin identity owners (individuals and organisations) who have earned the level of trust necessary to register new self-sovereign identities on Sovrin are called "trust anchors". Trust anchors play a special role in the Sovrin Trust Framework. They can also vouch for new individuals and organisations to become trust anchors. In doing so, they will enhance their own reputation.

Sovrin stewards, the engines of the Sovrin network, have much to gain from early participation. For more about Sovrin stewards, see the Sovrin Infrastructure Providers section of this paper and the white paper "Becoming a Sovrin Steward".

Many stewards will also be issuers of verifiable claims, as well as being consumers of claims issued by others. To many, the claims from stewards will be more highly regarded than those from organisations that are less well trusted, and thus “worth” more. Individuals will be more likely to want the claims of such an organisation in their Sovrin identity, further enhancing the steward’s reputation and stoking the virtuous cycle.

Sovrin will provide a mechanism for identity owners to publicly or privately share their “reputation graphs”, the integrity of which can be verified cryptographically via the Sovrin ledger. Unlike models based on self-attested social media activity, Sovrin reputation graphs will be based on a mixture of verifiable attributes and claims from trusted organisations and individuals as well as self-attested data. This will give Sovrin identity owners a powerful new tool in the emerging reputation economy.

Incentives for Stewards

Stewards who operate Sovrin nodes will be able to perform reads and writes more efficiently than others simply because they have direct access. All others need to request read/write access via a steward’s nodes. This lower latency, first place queueing and controlled service quality provides the first incentive for an organisation to serve as a Sovrin steward.

Incentives for Readers and Writers

Rather than spending vast amounts of money accumulating personal data from many different, expensive silos, Sovrin will enable organisations to simply ask their customers or partners to provide that data directly, instantly, verifiably and digitally from their Sovrin identities.

These “readers” of Sovrin identity data are termed “relying parties”. The improvements in access, security and privacy are obvious benefits to relying parties. But this data will only become available for identity owners to present to relying parties if it is first placed there by a “writer” of the data, called an “issuer”. In short, the need for both readers and writers is classic two-sided market.

Being an issuer of claims such that the customer’s relationship with an organisation transcends any single interaction is a powerful incentive, and by far the most customer-centric approach to managing the myriad identity data which must be accrued in the regular course of business. Many relying parties will also be issuers, serving to kick-start the ecosystem.

New Global Marketplace for Premium Claims

By establishing a common identity layer and creating a level playing field for all parties in the identity ecosystem, Sovrin enables the creation of a new global market for verified identity attributes, also known as “verifiable claims”.

It is likely that some issuers of verifiable claims will want to see a return on the investment in creating those claims, such as the cost of verifying a physical document. To do this, the Sovrin Trust Framework will define how issuers can set a price for the claims that they write to an individual's Sovrin identity. A relying party who wants access to these claims must agree to pay this price in order to use them. Verifiable claims that carry a price tag are called "premium claims". Premium claims are not part of the initial release of Sovrin, but will be introduced as Sovrin evolves and gains further traction.

It is not envisaged that government-issued public credentials such as birth certificates, driving licenses, marriage certificates etc. will be premium claims (though there is nothing to stop that). Schools, universities and other public institutions are also unlikely to charge for premium claims as it is part of their charter to provide these attestations on behalf of their members.

However private sector issuers will have a choice: provide verifiable claims for free or charge for them. It is entirely a free market decision.

Note that individuals may also issue their own premium claims.

When a relying party pays for a premium claim, a proportion of that payment will be shared with the Sovrin Foundation (the "brokering fee"), and the remainder will go to the issuer. This brokering fee may be split three ways (Fig 7):

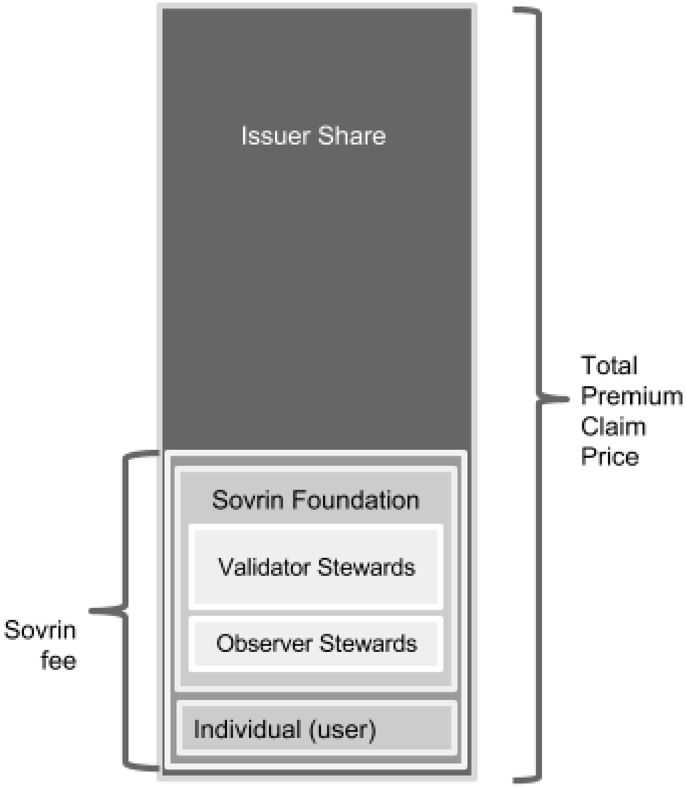

Fig 7: Premium Claims revenue split
Note: proportions are not to scale

- A portion for the operational overhead of the Sovrin Foundation itself.
- A portion for the costs to stewards of running Sovrin nodes.
- A portion to the individual identity owner who is the ultimate source of the claims.

Using this model, governance work of the Sovrin foundation can be funded directly from the revenue from the marketplace for premium claims enabled by the Sovrin identity Network. This is a sustainable long-term funding model for effective governance.

Sovrin Infrastructure Providers

The Sovrin Identity Network is a global public utility. On top of it will exist new products and services which are created in a competitive open market (Fig 8).

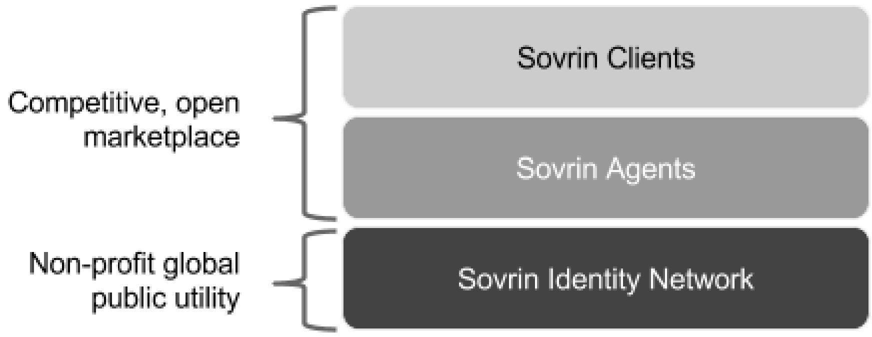
Competitive, open marketplace
Non-profit global public utility
Sovrin Clients
Sovrin Agents
Sovrin Identity Network

Fig 8: the layered approach of Sovrin

In the words of Smith and Khovratovich[21] in their paper "Identity System Essentials", "the identity system is the hat rack for transactions that are the supported hats". Sovrin is the hat rack, and an open marketplace will create the hats.

The major roles that are intrinsic to Sovrin infrastructure are described in the following section. Note: it may be useful to refer to the Sovrin Glossary and the Technical Foundations of Sovrin white paper.

Stewards

A Sovrin steward is an organization that has permission to operate a node on the Sovrin network. More information on participating as a Steward can be found in the white paper "Becoming a Sovrin Steward". A steward:

- Has the will and capacity to operate a validator node on the Sovrin network per the published specifications.
- Agrees to abide by security, privacy, reliability, and other policies of the Sovrin Trust Framework.
- Believes their organization can benefit from a secure and independent source of verified credentials, claims, and reputational data for their customers, clients, partners, vendors, or members (constituents).
- Wants to help to shape the future of Internet identity.
- Wants to strengthen their image/reputation in the eyes of their constituents.

Stewards will be organisations in a position of public trust and/or which have a vested interest in the integrity and success of Sovrin, such as healthcare providers, NGOs, financial institutions, universities, governments, and dedicated Sovrin service providers.

As the only organisations entitled to run Sovrin validator and observer nodes (see "The Technical Foundations of Sovrin"), stewards benefit from the lowest latency and highest performance for accessing Sovrin (likely to be particularly important for high volume users).

The first stewards, called founding stewards, will be selected by the Sovrin Foundation Board of Trustees from among the early participants in Sovrin infrastructure. Founding

[21] Identity System Essentials 2016: Samuel Smith Ph.D. and Dmitry Khovratovich Ph.D.

stewards will have significant influence on the shape of Sovrin; will be involved in the creation of the Sovrin Trust Framework; and will help nominate and elect additional stewards.

Trust Anchors

Organisations who have permission to create new identities on the Sovrin ledger are known as “trust anchors”. A trust anchor must have some pre-existing knowledge of the identity owner for which the identity is being created. As such, they are also likely to be issuers of verifiable claims which the identity owner can utilise elsewhere.

An example of a trust anchor could be a credit union that is legally required to “know your customer” (KYC) and already has a secure way of identifying them online. Such an organisation would be very well placed to register a Sovrin identity for a customer, thus making the customer the identity owner. Once established on Sovrin, the identity owner can begin accepting claims from other issuers to their Sovrin identity. (Note that only the identity owner can tie these claims tie together—there is no external correlation between them unless the identity owner wishes to reveal it, for example by sharing claims from two different issuers with the same relying party.)

Trust anchors thus perform a vital service within Sovrin—they are the means by which new identity owners are brought on to the network while maintaining the network’s overall integrity and trust. (Note that all stewards are automatically trust anchors as stewards are in the highest level of trust on the network.)

Agents & Agencies

An agent is a software process acting on behalf of an identity owner using one or more of the identity owner’s Sovrin identities. In most cases, an agent will be hosted by an agency (i.e. a service provider specialising in hosting Sovrin agents), although an agent can always be self-hosted by an identity owner.

An agent provides a persistent addressable endpoint by which an identity owner can send and receive notifications, messages and other transactions to and from other identity owners. Agencies can also offer anonymised endpoint services for maintaining privacy and accountable pseudonymity.

Another service agents and agencies provide is permissioned sharing and secure storage of off-ledger data. This can include audit logs, large files or other activities which may be too numerous or ancillary to store in the Sovrin ledger, such as authentication attempts.

Sovrin agents can be developed and hosted by anyone, just like an email server or a web server. We anticipate a healthy new market in innovative Sovrin agent technology will develop. Identity owners will be able to freely move between agents to take advantage of new features, encouraging real competition.

It is envisaged that artificially intelligent (AI) agents will be created that act on the identity owner's behalf, following rules laid down by the identity owner. These AI agents will simplify peoples' online lives while still ensuring that they retain control, security and portability.

For example, an AI agent may know when a person's car insurance is due to expire, and automatically request quotations from providers using minimal, yet verifiable details about the person. The AI agent may then act on the person's preferences and rules to select the best policy, execute the transaction, and inform the person that this has been done.

Issuers & Claims

Issuers are people, organisations or things that issue "claims" (attributes) about a Sovrin identity owner. A claim is an identity record that asserts one or more attributes of a Sovrin identity, such as the fact that you have a driving license, you live at a certain address, or you like a particular type of food. Issuers may be organisations such as driving license providers, banks, energy companies, insurers and so on.

Claims can be also self-asserted. For example an identity owner can issue a claim that they want to move house to a certain area under a certain budget, and can then seek interest from real estate agents, moving services, and mortgage providers.

Verifiable claims are those which have been issued by party other than the identity owner. For example, the passport office may issue a verifiable claim that a person has a passport with a specific number and expiration date. To be verifiable, this claim would be signed by the passport office using its own Sovrin identity and the indelible proof of the claim's existence (a hash) would be stored on Sovrin, such that any relying party can be certain of the claim's origin.

An identity owner's Sovrin identity is therefore made up of a number of claims created by a number of issuers. Together these form a set of Sovrin identity records. Identity owners may have as many identity records as needed, which can be combined into any number of personas in order to protect the identity owner's privacy.

Issuers are able to revoke claims (or individual elements of claims) which they have issued. For example, a person may be banned from driving for a year, so their entitlement to drive is revoked. However, the name, address and date of birth on their driving license remain valid.

Issuers are highly likely to be reliant on claims created by other issuers to carry out their business. Therefore a contributory network effect will result in more issuers wanting to write more claims because they also benefit from the claims written by others.

Relying Parties

A relying party accepts claims from a Sovrin identity. For example, a person wanting to set up a new energy contract might use Sovrin to share a proof of address with energy companies. Those energy companies are relying parties.

A relying party may directly request specific claims from a Sovrin identity owner, such as a name, address, and date of birth. The identity owner (or their agent) will be able to select the claims from their list of Sovrin identity records which match the request. The identity owner can then consent to share the data with the relying party, and the resulting consent record can be written to the ledger so both parties can prove consent was given.

The relying party can verify signatures on claims from different issuers by looking up the issuer's own Sovrin identity to obtain their public key. The signature also verifies that the data has not changed since it was first written. This enables the relying party to instantly confirm details about the identity owner that can take hours, days, or weeks to verify via conventional channels.

Core to the design of Sovrin are the concepts of non-correlation and anonymity, which are described in more detail in The Technical Foundations of Sovrin. In short, using Sovrin, it is possible for the identity owner to prove their identity by satisfying certain properties requested by a relying party in an uncorrelated way without revealing other identity details.

The Road Ahead

Sovrin has been developed explicitly to deliver the Internet's missing identity layer. Using distributed ledger technology it has become possible for the first time to create a global "identity operating system" which puts the individual in control, provides high security, and also delivers portability.

Self-sovereign identity will improve the way Internet services work. Reliance on outdated, silo-based identity models will fall away, along with the high costs of maintaining them. Existing services will evolve rapidly to take advantage this new public utility for identity, and new services and markets will develop which could never have existed before.

But this won't happen instantly. Sovrin relies on a two-sided market of issuers and relying parties cooperating and recruiting others. As the Internet age has progressed, the speed with which two-sided markets can emerge has accelerated hugely. For example, in March 1876 there were only two telephones in the world, so it took over a century for telephony to become truly universal. By comparison, WeChat has grown to 700M monthly active users in just over two years.

As a new two-sided market, Sovrin's growth strategy is threefold:

1) Establish and grow the Sovrin Foundation's membership, create and evolve the Sovrin Trust Framework, and establish sustainable, reliable governance of the Sovrin ledger.
2) Grow the number of issuers, and the relying parties will come (it's helpful that most issuers will also be relying parties).

3) Open up Sovrin to the global developer community to innovate and create new features and new markets that were previously impossible.

The initial developers and organisations using Sovrin will create islands of functionality around specific use cases. Some will be small, some will involve millions of people, and some might not even leverage self-sovereignty as a headline feature. As more organisations come on board, more islands will be created, grow in size, and begin to overlap. As this happens, individuals will realise that the Sovrin identity claims they are using with one organisation can be used seamlessly with another, and another, and another, with complete control and confidence, thus unlocking the network's full potential.

The most forward-thinking organisations will already realise how beneficial it will be for them to be recognised as Sovrin trust anchors that give their customers the ability to use their identity in other places, rather than "capturing" or "farming" them. They will also see Sovrin's potential to substantially improve their customers' digital experiences, stripping away decades of inefficiency, substandard security, and lack of privacy. Innovation will blossom.

We at the Sovrin Foundation are excited about the road ahead. Sovrin is the identity layer the Internet has been waiting for, and its emergence will be as transformational as the Internet itself.