

Sovrin: What Goes on the Ledger?

A white paper from Evernym in cooperation with the Sovrin Foundation.
An overview of what's on the Sovrin distributed ledger and why.

The logo for Evernym, featuring the word "evernym" in a lowercase, sans-serif font. The "e" is green, and the "er" is grey. A green dot is positioned above the "e", and a green line connects the "e" to the "r".The logo for Sovrin, featuring a stylized orange sunburst icon to the left of the word "sovrin" in a lowercase, sans-serif font. Below "sovrin" is the tagline "identity for all" in a smaller, lowercase, sans-serif font.

Andrew Tobin, Evernym
April 2017

evernym.com
sovrin.org

Executive Summary

One of the most common questions the Evernym team is asked when discussing Sovrin is, “What’s on the ledger?”. This paper answers that question, giving the background behind some of the decisions that have been made about what data goes on the ledger and what must not.

Sovrin’s focus on decentralised, independent identity, has resulted in policies that are privacy enhancing and privacy protecting. Every decision made in its design has been focused on increasing security and privacy while reducing undesirable impacts and correlation risk.

The Sovrin ledger is designed specifically to be a public utility which can be used by any person or organisation. This has huge benefits from a scale, flexibility and adoption perspective. This public accessibility ensures that identity can be truly independent and self-sovereign.

As a result, some of these policies will differ substantially from what is seen in other ledger implementations where identity is a “bolt on” rather than the sole purpose of the ledger itself, as is the case with Sovrin.

Background

Distributed ledgers are increasingly being seen as a way to deliver independent identity. Independent identity is a privacy-protecting digital existence for people, organisations and things that is separate from any data silo. Independent identity is owned, managed and controlled directly by “Identity Owners” rather than third parties.

The Sovrin identity network, launched by the Sovrin Foundation in September 2016, is the world’s first distributed ledger engineered specifically for independent identity. Further detailed documentation about Sovrin can be found in the [Sovrin library](#).

Independent identity is the solution to the problem caused by a fundamental flaw in the design of the Internet: because the Internet was built without an identity layer, every service must build its own data store (“silo”) to store and secure customer details. Huge silos have amassed data for billions of users, each one able to monitor, suspend, delete and remove any of its users at will.

This approach has led to security- and privacy-damaging results, as seen by the ever increasing scale of hacking and breaches, tracking, and un-permissioned data sharing that is so commonplace in today’s Internet.

With the emergence of distributed ledgers, it has become possible to create a public global ecosystem where an individual can own their digital identity, rather than “rent” tens or hundreds of identities from different silos. With distributed ledger technology, anyone in the world can have a self-sovereign digital existence.

Independent identity brings benefits that will transform the way online services are developed. [Sovrin](#), developed by [Evernym](#) and donated to the Sovrin Foundation, who then released it as open-source, is the world’s first independent identity network.

Identifiers, Keys and Endpoints

At its heart, Sovrin is a registry for Decentralised Identifiers¹ (DIDs) and their associated verification keys.

Decentralised Identifiers Explained

DIDs are a type of identifier designed for verifiable digital identity that is “independent” or “self-sovereign”, i.e, fully under the control of the Identity Owner and not dependent on a centralised registry, identity provider, or certificate authority. Developed for distributed ledger applications, DIDs are the foundation for decentralised identity management.

Each DID stored on Sovrin resolves to a DID descriptor object (DDO), also stored on Sovrin, that contains all the metadata needed to prove ownership and control of a DID as well as share the cryptographic keys and resource pointers (or endpoints) necessary to initiate trusted peer interactions between Sovrin entities. Sovrin determines how a DID is registered, resolved, updated, and revoked.

Due to its tamper-resistant nature and verifiable transaction chronology, users can be reliant on Sovrin to return the current verification key for any DID. A DDO contains the current verification key that the Identity Owner of it’s DID is using, and can be signed by the private key of the Identity Owner, thereby establishing proof of ownership of the DID. This solves one of the biggest issues with existing public key implementations: knowing which key is currently in use.

Agent Endpoints

A DDO also contains a pointer to an “endpoint”. The endpoint is the location the Identity Owner wants to use for further communication, such as a URL or IP address. It enables two DID owners to communicate securely directly with one another in a private, trusted peer-to-peer interaction, with no intermediaries.

¹ Decentralised Identifier specification:
<https://github.com/WebOfTrustInfo/rebooting-the-web-of-trust-fall2016/blob/master/draft-documents/DIDSpecificationWorkingDraft04.pdf>

An endpoint will typically point to an Agent. An Agent is a software program or process acting on behalf of an Identity Owner to facilitate interactions with other Agents or the Sovrin Ledger. Agents provide persistently addressable communication endpoints to, for example, store communication requests while the Identity Owner is offline.

Enhancing Privacy With Pairwise DIDs

To protect privacy and avoid correlation, each Identity Owner can have multiple DIDs stored on Sovrin, and by definition this means they can have multiple verification keys and endpoints.

Using a single unique identifier for all transactions can be very damaging for an Identity Owner, as it allows them to be correlated across all those transactions. In non-Sovrin distributed ledger applications, where such transactions are written permanently to the ledger, this approach results in an indelible record of all the interactions of an Identity Owner which is visible to all who access the ledger

With Sovrin, each user is able to establish a specific DID for every relationship that they have. These are called “pairwise” identifiers and they are tremendously powerful. They allow an Identity Owner to keep their interactions with one party completely separate from any other party. The Identity Owner is the only person able to map and correlate their DIDs. This means that the relationship a person establishes with their bank is entirely separate to the one that they establish with a retailer, so neither bank nor retailer can correlate the Identity Owner’s activities with each of them.

Sovrin users can also have a DIDs and associated DDOs containing publishable information. An example would be a bank publishing a common DID which any potential customer can use to get in touch with them and establish a relationship. This DDO might contain details about the bank, such as their name, address, opening hours, and public-facing endpoint.

Example of DID Use

This section explains at a high level how DIDs are used on Sovrin.

Independent Identity Owner Andy (using Sovrin) wants to apply for a loan with Seagull Bank. Seagull Bank publishes their DID on their website and all their literature and advertising using a QR code. Andy sees a particularly attractive offer on a poster in a bus shelter, and scans the QR code with his phone.

Andy’s Sovrin-enabled app recognises the QR code as a Sovrin DID. It accesses Sovrin to look up the DID, and access the current DDO. It verifies that the DDO is valid and signed by the owner of the DID. Within the DDO is the pointer to the bank’s endpoint.

Andy establishes² a new DID for this relationship with Seagull Bank, with associated DDO containing a new public verification key, plus the endpoint that he wants Seagull Bank to use. Andy then contacts the bank directly, peer-to-peer with no intermediaries, using the bank’s endpoint, and provides the bank with this DID.

The bank now has a way of communicating directly and securely with Andy via the endpoint he has provided in his DDO, and does so, offering him specific details on their loan offer. Andy can use the DID & DDO information previously retrieved from Sovrin to verify and decrypt the communications from the bank, and when he responds to the bank the bank can do the same.

Thus Andy and Seagull Bank can use Sovrin to find each other, and once found they can use the DID & DDO information returned to establish a trusted peer interaction with no intermediaries, which is totally private and unique to that particular relationship. They can then share verifiable claims or any other information through that secure channel. This is like Andy calling the bank and instantly inventing a brand-new language together which only Andy and the bank knew. If someone were to tap the line and listen, they wouldn’t understand a thing.

In this example, Sovrin is only used for the lookup and discovery of DIDs. Once the DIDs of the two parties are established, all further communication and data exchange takes place off-ledger. Sovrin is used to verify the parties involved, and confirm the validity of verifiable claims and proofs which may be then exchanged completely off-ledger.

Schemas and Claim Definitions

Schemas and claim definitions are stored on Sovrin.

A schema specifies the data types and formats which are used to make up claims.

A claim definition is published by a claim issuer (e.g. a bank, passport office etc). It references the relevant schema, the issuer that published the claim definition, and the signature types used.

Claims contain data about an Identity Owner. Claims can be issued by 3rd parties or can be self-attested. Identity owners using Sovrin accumulate claims into an identity account or vault (usually held off-ledger - see next section) and use them to prove who they are.

Verifiable claims³ are the currency of independent identity. They are defined thus:

Verifiable claim (aka third party claim) is a claim issued by a third party, containing data about an Identity Owner, and signed by the issuer.

² To establish a new DID, the Identity Owner creates the DID, DDO and associated key pair, and requests one of their trust anchors to write this to Sovrin.

³ W3C Verifiable Claims Task Force definition: <https://w3c.github.io/vctf/>

ç^!ãã\^Á Á@Á^}•^Á@eÁ Á!ãã Á eÁ^Á^!ããáÁ^ Á Áããã} æ!^Á} Á@Á
Ù[ç!ã Á^} æ Á^& !ãÄ
Á

Claims can also be self-asserted. These are claim records asserted by the Identity Owner whose identity record it describes. For example, Andy may claim that he is a fan of Manchester United, on his own authority, and Manchester United does not have any involvement.

In order to make sense of the content of a claim, a claim definition is used. For example, a driving licence authority wants to issue people with a digital version of their driving licence in the form of a verifiable claim. The first thing that they do is publish the claim definition onto Sovrin. This definition contains the structure of the claim they are going to issue including the references to attribute names and types from a relevant schema, such as the driver name, address, driving licence number, date of issuance, vehicle types permitted etc.

Enabling claim issuers to publish their own claim definitions has significant advantages.

It enables huge scale and encourages agility. Because no central authority is needed, the restriction of approval for issuance of new claim types is removed. This lets the smallest organisation, or a specific individual create a new claim type which can be instantly accessed by any other Sovrin user in the world.

The restrictive and cumbersome hub-and-spoke model of traditional attribute and data exchange is replaced by a flat layer, where claim definitions drive agility and flexibility and there are no barriers to entry for creating new claim types. We feel this approach will best foster the organic development of widely used, cross-industry definitions.

Claim definitions are also searchable. A Sovrin user may wish to find out who issues claims containing certain data types. For example, finding which banks issue claims containing a proof of banking relationship (which can be very useful when renting a house for example). The person may then elect to set up a bank account with one of those banks, rather than one that does not issue Sovrin claims, in order to be able to obtain such a claim for use in their digital life.

No Private Information on Sovrin

Evernym and the Sovrin Foundation does not recommend storing private personally identifiable information on Sovrin. Indeed Evernym strongly advises against any storage of these information types on any ledger.

Sovrin contains a permanent record of transactions. Storing private information on a distributed ledger creates an indelible record of that event. This is not privacy enhancing.

Should there be any compromise of your keys, or if the relying party or issuer is compromised, an attacker may be able to retrieve this record of your data which you can

never deny existed. As well as being personally concerning, this also affects user rights under new legislation such as the European General Data Protection Regulation.

Evernym’s policy therefore is to keep private data, including hashes of private data, off the Sovrin ledger.

Proof of Sharing Consent

When one Sovrin Identity Owner shares data with another, it may be desirable for both parties to record that fact. For example, a retailer may need to prove to a privacy auditor that they had the permission of a customer to receive and store that customer’s data.

As discussed in the previous section, it is not desirable to store private information on Sovrin. Therefore how can sharing consent be handled in a way that allows both parties to prove that information was shared between them?

The answer is for the two parties to each store a consent receipt privately, and record the proof of that consent receipt, containing no private information at all, on Sovrin. This is a proof of existence of the sharing agreement which, thanks to the immutability of the ledger, will allow either party to provide irrefutable, auditable proof of the terms if required in the future. The consent receipt contains the two DIDs that shared the information, and the attribute names and data types that were shared (rather than the actual data itself), and is signed by both parties to provide non-repudiation. The proof which is stored on Sovrin is simply a cryptographic hash of this receipt.

Because hashed data stored on Sovrin is opaque, only the parties who already know the substance of the agreement will be able to disclose its existence, and since DIDs are not readily correlatable back to an individual Identity Owner, it would be prohibitively difficult to “mine” the ledger for evidence of agreements between known parties.

Public Claims

While Evernym’s policy is to keep private data off any type of ledger, there are many situations where public data may be stored on Sovrin as a public claim or as part of a DDO.

For example, a company or organisation may elect to publish digitally verifiable versions of data that is already part of the public record. For example the company’s registration information, the list of directors, the company’s address. The company could also publish verifiable claims from regulatory authorities that confirm the company’s licence to trade particular goods for example.

Sovrin can therefore be used as a repository for public claims, as long as the issuer has a clear understanding that what they are publishing is a permanent record of that data which cannot be deleted.

Revocation Registries

Revocation is a vital part of any identity ecosystem. A common question that is asked about identity on distributed ledgers is “how can a credential that is permanently written to a ledger be revoked?”

As stated in the previous sections, private credentials or claims are not written to Sovrin in the first place. They are supplied by an issuer to an Identity Owner in the form of verifiable claims. Therefore there is still a requirement to revoke a verifiable claim, or a single attribute within one.

Much like attribute exchange and claim verification, most traditional approaches to revocation require the relying party to connect directly with the issuer or some other central authority to check validity of the provided data. This presents a privacy risk, as issuers and relying parties can correlate an Identity Owner across domains.

An always-accessible list of which credentials have been revoked needs to be made available, meaning that if the issuer’s system is down no revocation checks can be carried out. It also places a large technical burden on the issuer, who needs to create and maintain an API, vet all those who want to connect to it, and on the relying party who will need to create and maintain many such APIs to every issuer around the world. Lastly, it creates complex and unwieldy user experiences.

For example, in order for a hire car company to check if a UK driver’s licence is valid, the driver must first, ahead of time, obtain a “check code” from the UK driving licence authority (the DVLA). To obtain this check code the driver must provide personal information including their national insurance number and postcode.

The DVLA warns of data “leakage” that occurs as a result of this transaction, stating “details from your DVLA record and your National Insurance number will be shared with other government departments (HMRC and DWP) to check your identity”.

The DVLA then provides the driver with a check code which is valid for 21 days, which the driver must then provide to the hire car company. The car hire company uses the check code as user authorisation to execute a query on the DVLA’s database. This complex interaction is necessary because the DVLA cannot permit unrestricted access to its database without driver consent.

In the process the DVLA is able to see and correlate a given driver’s activities, which may be undesirable to the Identity Owner.

With Sovrin, revocation is decentralized, asynchronous, and private. The issuer publishes a revocation registry to Sovrin at whatever frequency is desired. This registry references the relevant claim definition, and contains a cryptographic accumulator used for proving and

verifying set membership. You can think of this accumulator as an exotic kind of compound hashing function - its value changes when credentials are added to or removed from the list, but it is impossible to directly test whether a particular credential is contained within it without being the credential owner..

Using their knowledge of which credential belongs to them, the Identity Owner is able to create a zero knowledge proof that their credential belongs to the set of valid credentials (without disclosing which one it is) at the same time as the rest of the desired attributes are proved.

The relying party can verify the proof using a recent copy of the accumulator (just like they might use a recent copy of someone's public key to verify a signature). Therefore the relying party can confirm if an attribute provided by an Identity Owner is revoked or not, without having to contact the issuer of that attribute.

This preserves privacy while maintaining the ability to revoke a credential at any point.

It also creates a completely new and highly efficient revocation mechanism which can transform the personal data economy. Because a relying party can check the validity of data provided to them without needing to contact the issuer, huge layers of complexity are removed from the current models of data exchange.

Summary

This paper has described Evernym's policies for the data that does and does not go on Sovrin, and the logic behind those policies.

Using Sovrin and the mechanisms described in this paper, a relying party can, for the first time, do all of the following without ever needing to contact the issuer of a claim:

- Confirm that the data provided to them by an Identity Owner did come from the stated issuer;
- That the data is unchanged;
- That the data was provided only to the Identity Owner who has presented it;
- That the data has not been revoked by the issuer;
- And that the Identity Owner consented to the sharing of the data.

For an issuer, the consequences are similarly profound and liberating. An issuer can:

- Create and issue any type of claim without waiting for a central body or data hub to update its limited transaction set to accommodate it;
- Revoke credentials or claims that it issues, without needing to create complex and privacy damaging user experiences or handle multiple technical integrations with thousands or millions of relying parties;

- Provide its customers or users with trustworthy digital credentials that can be used anywhere in the world instantly.

And all of the above can be done in a highly privacy preserving manner, protecting the Identity Owner from intended or unintended correlation and inadvertent data leakage.

Taken together this represents a revolutionary way to manage and share personal data. Old-style hub-and-spoke attribute exchanges are no longer required.

In summary:

Goes on Sovrin

- Decentralised identifiers and associated DDOs with verification keys and endpoints.
- Schemas and claim definitions
- Proof of consent for data sharing
- Public claims
- Revocation registries

Does not go on Sovrin

- Private data of any kind (including hashed personal data)
- Private proof of existence