How Sovrin Works

A Technical Guide from the Sovrin Foundation

by Phillip J. Windley, Chair, Sovrin Foundation

3rd October 2016

sovrin.org

Introduction

Sovrin is an open-source identity network built on distributed ledger technology. Sovrin is public and permissioned. Public means everyone can use it. Permissioned means that the network nodes that ensure consensus of transactions on the ledger are governed, in this case by the non-profit Sovrin Foundation.

In this discussion we're going to look at the interactions of a Sovrin user named Jane with her bank, her local government, a potential employer, her school and a retailer.

The following figure shows Jane's view of her identity on Sovrin. Right now there's nothing there, but we're going to add things as we discuss Sovrin's capabilities in the following sections. Jane's identity doesn't really exist as depicted. The view is a virtual representation. Jane's Sovrin identity is the collection of all of her Sovrin identifiers, claims, disclosures, and proofs. The things in the box labeled "Jane's identity" are stored in various places. Most, but not all will be on the Sovrin ledger itself, some might be stored off the ledger in other repositories like the private ledger we'll discuss later.

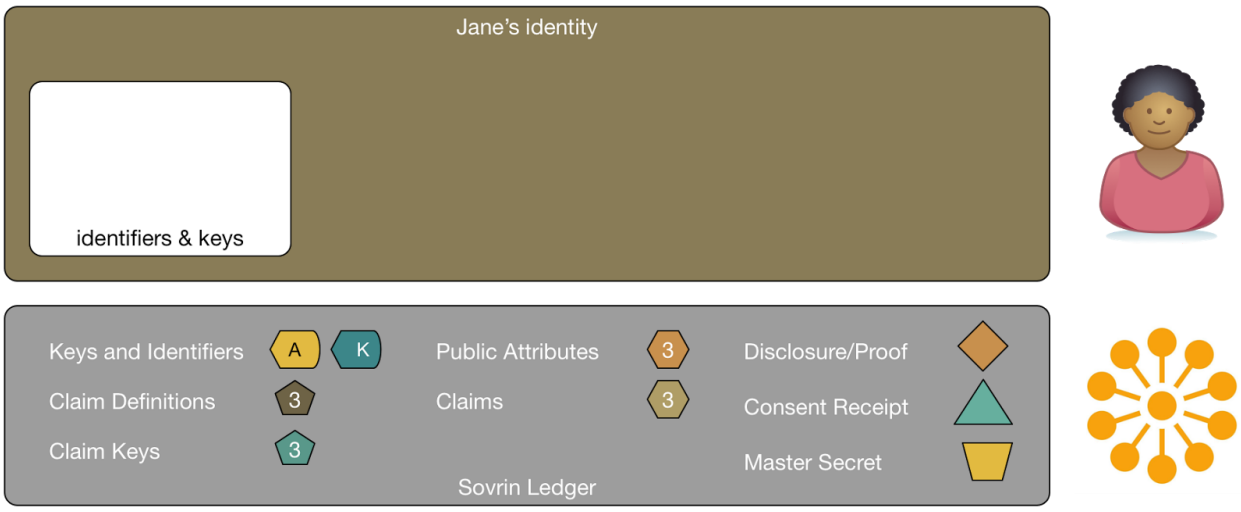

Jane and an empty ledger

The diagram also shows the Sovrin Ledger. The ledger is shown to emphasize that everything we talk about is using the ledger. There are far too many lines for the diagram to show all the various interactions with the ledger itself, so I've chosen to merely represent it and use it as a place to show the diagram's legend.

The Sovrin Identity Network (SIDN) consists of multiple, distributed nodes located around the world. Each has a copy of the ledger. Nodes are hosted and administered by *stewards*. Each node has a copy of the ledger. Stewards are responsible for validating identity transactions to assure consistency about what is written on the ledger and in what order. They do this using a combination of cryptography and an advanced Byzantine fault tolerance algorithm. See the Sovrin whitepapers The Inevitable Rise of Self-Sovereign Identity (PDF) and The Technical

Foundation of Sovrin (PDF) for more details.

Keys, Identifiers, and Relationships

Sovrin tracks keys and identifiers. One of the major concerns with identity is *correlation*. If Jane were to use one identifier in multiple places, those places might collude to correlate that identifier and amass significant data about her without her permission. Sovrin avoids this by allowing Jane to use a different identifier with everyone she relates to.

By default, Sovrin identifiers are cryptonyms, an encoded Ed25519 digital signature verification key. Sovrin also supports DID's (Distributed Identifiers), which are identifiers with no cryptographic properties. These identifiers also have an associated Ed25519 verification key. In the diagram the signing key is represented by a small letter and the verification key is represented by a big letter. These two keys represent a private-public key pair. Jane never shares her signing key, only the verification key.

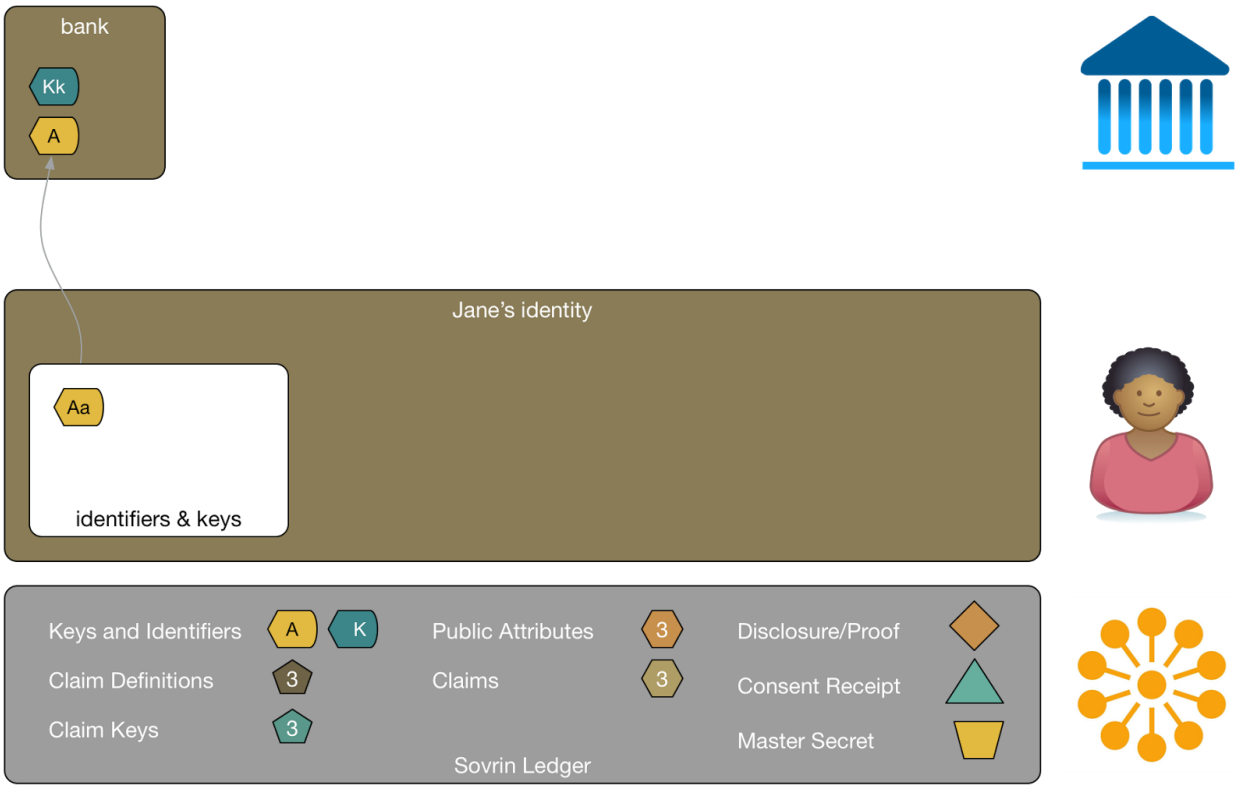

Jane's relationships

Jane has a relationship with her bank. She shares a verification key, A, with the bank that created specifically for this relationship. This key represents Jane's identity to the bank and can be used to verify any interactions that they have. The bank also has its own key, K. This is a well-known key that represents the bank to the world. Jane would also have a copy of that so that she can validate communications she has with her bank. The verification keys of both Jane and the bank are found on Sovrin, so they can both know they are using the latest verification

keys of the other party.

As we add new relationships to this diagram, you'll see that Jane uses a unique key pair for each of her relationships.

Claims

In addition to identifiers, Jane has claims on Sovrin. Claims are a fundamental component of Sovrin. They are assertions or attestations made by a party about itself or another. Claims are digitally signed so that anyone receiving the claim can know who issued it.

There are several types of claims. In the following diagram, claim 1 and 4 are self-asserted. Jane might, for example, create a claim that asserts her gender or that her name is Jane.

The other claims are verifiable claims made by others about Jane. Claim 2 was asserted (and signed) by the local government. This claim might be a driver's license. Jane could use this verifiable claim to prove to someone else that she's authorized to drive.

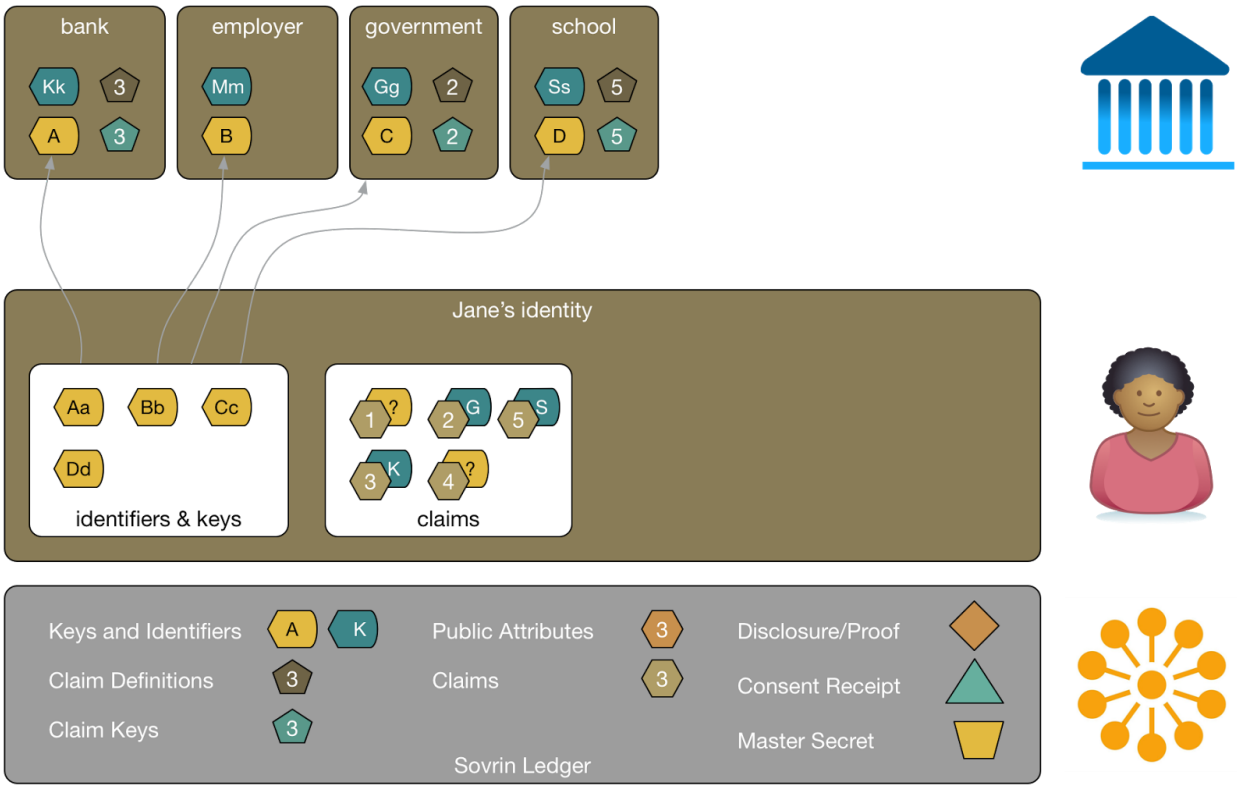

Claims about Jane

A claim in Sovrin is specific to an identity owner, its subject. Claims are linked to the identity owner and the party issuing the claim. For example, claim 2, Jane's driver's license, can be validated as being about Jane and as having been issued by her local government.

Claims are defined so that they are understandable by parties that rely on them. Sovrin makes provision for claim types to be defined with a schema or ontology. Claim definitions are recorded

on the Sovrin ledger. The local government would create a claim definition that describes a driver's license claim. Jane's claim includes a reference to the definition so that anyone to whom Jane presents her driver's license can look up its definition and make sense of it.

Claims also have keys that are specific to the claim so that relying parties can validate that the claim is legitimate. Claim keys are linked to the issuer. A claim key can be changed, but old keys are saved, so that a relying party can verify a claim based on the keys that were used to issue it.

And, of course, claims can expire or be revoked when needed.

Disclosures

Claims can be reused and tailored to the applied purpose at hand using disclosure proofs. Disclosure proofs allow claims to be used without disclosing unnecessary information about the subject.

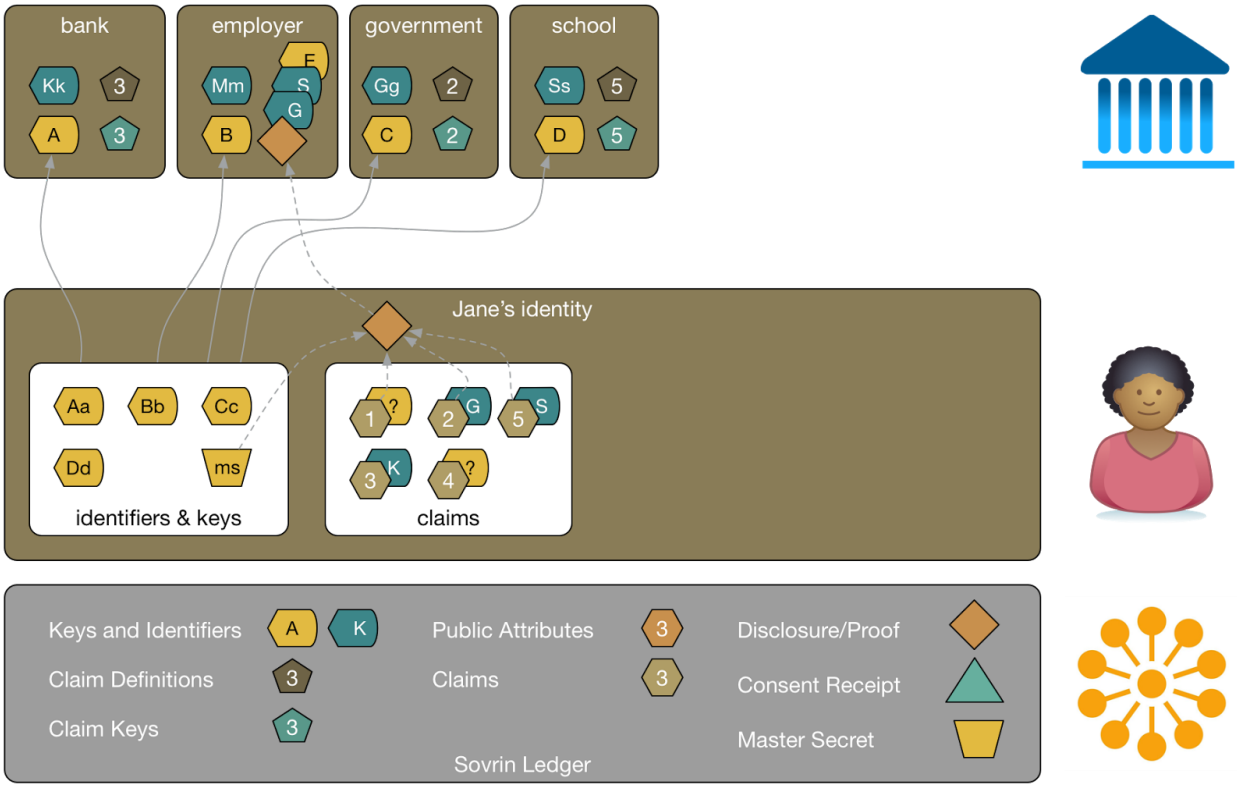

Jane discloses attributes using her claims

In this diagram, Jane is applying for a job. She can combine the verifiable claims from her school and government along with information she's self-asserted into a disclosure. Jane can pick the attributes she wants to share from each of these without disclosing the entire claim. For example, she might choose to include a proof from the government of her address and that she's over 18, a proof that she has a certain degree from her school, and a proof of her gender

from herself. Jane uses a master secret, a special key, to create the disclosure using a zero-knowledge proof algorithm. The disclosure proof links the attributes so that the employer knows that they were made about the same person.

The employer can verify each of these attributes were asserted by the party who makes the claim. The employer has a verification key from Jane that allows them to validate claims that Jane makes herself. Jane can disclose additional details as the relationship with the employer progresses. While the algorithms behind this are complex, the user experience is simple and natural.

An important point about disclosure proofs is that they are non-correlatable. Suppose someone at the employer had a friend who worked in the local government. These two people couldn't collude to use the proof Jane sent to the employer to discover additional information about Jane that the government holds. Neither the proof nor the identifier the employer holds for Jane correlate to any identifier that the government holds and thus can't be used as a key to look Jane up in the government's systems.

Claims Can Be Made About Anybody

Jane has a relationship with a retailer (we can tell since they have her verification key). Suppose that Jane wants some evidence that the retailer is in good standing with the local chamber of commerce. In the same way that Jane created a disclosure for a potential employer from claims, the retailer can use a verifiable claim from the chamber of commerce to create a disclosure proof and send it to Jane. She now has proof of what the chamber of commerce said about the retailer.

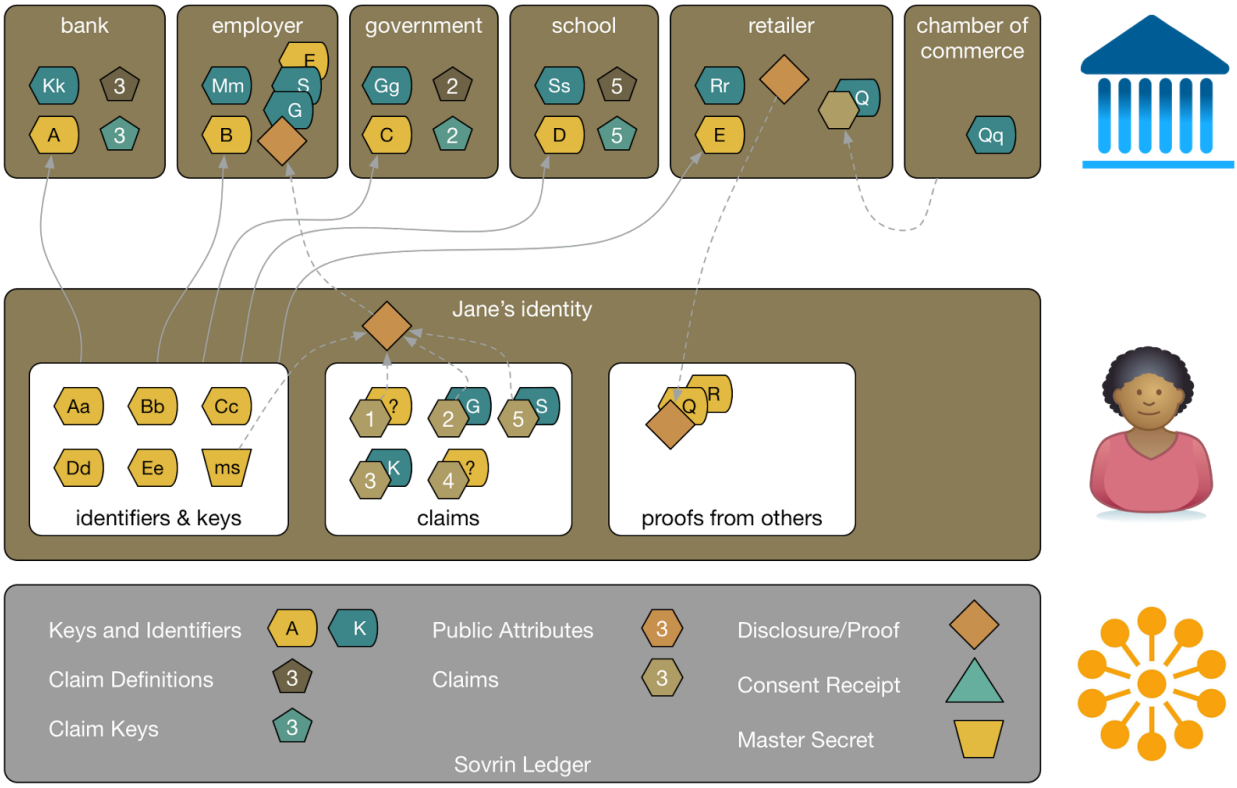

Jane receives a proof from a retailer

Consent Receipts

As part of buying from the retailer, Jane has shared personal information. The retailer can provide an attestation that they will treat Jane's personal data in a specific way (e.g permission to use expires after 30 days, attributes used only for completing transactions, and so on). That attestation is called a *consent receipt*. Jane uses Sovrin to keep track of the consent receipts she's received. Consent receipts are just one type of agreement that Jane can manage with Sovrin.

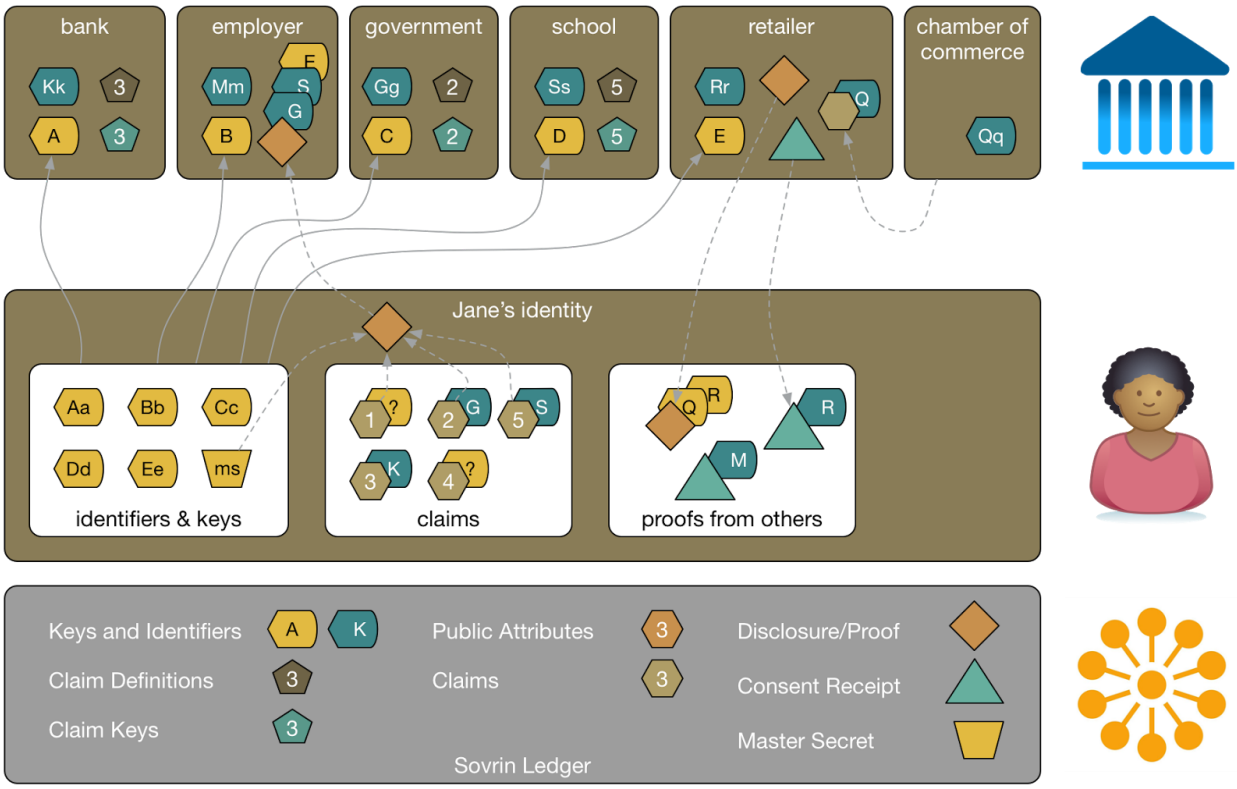

Jane receives a consent receipt

Public and Private Attributes

Claims can be public or private. For example, the school might have a public attribute showing its address. That isn't sensitive information, so it can be stored publicly on the Sovrin ledger. Anyone can read it and validate that it really is the self-asserted address of the school. Furthermore, others can attest to the validity of the claim, leading, over time, to increased trust that the address is legitimate. This kind of social proof is critical to establishing reputation and building trust.

Encrypted attributes are not normally stored on the Sovrin ledger. The following diagram shows that Jane has a private ledger. Things like proofs from others and encrypted attributes can be stored there.

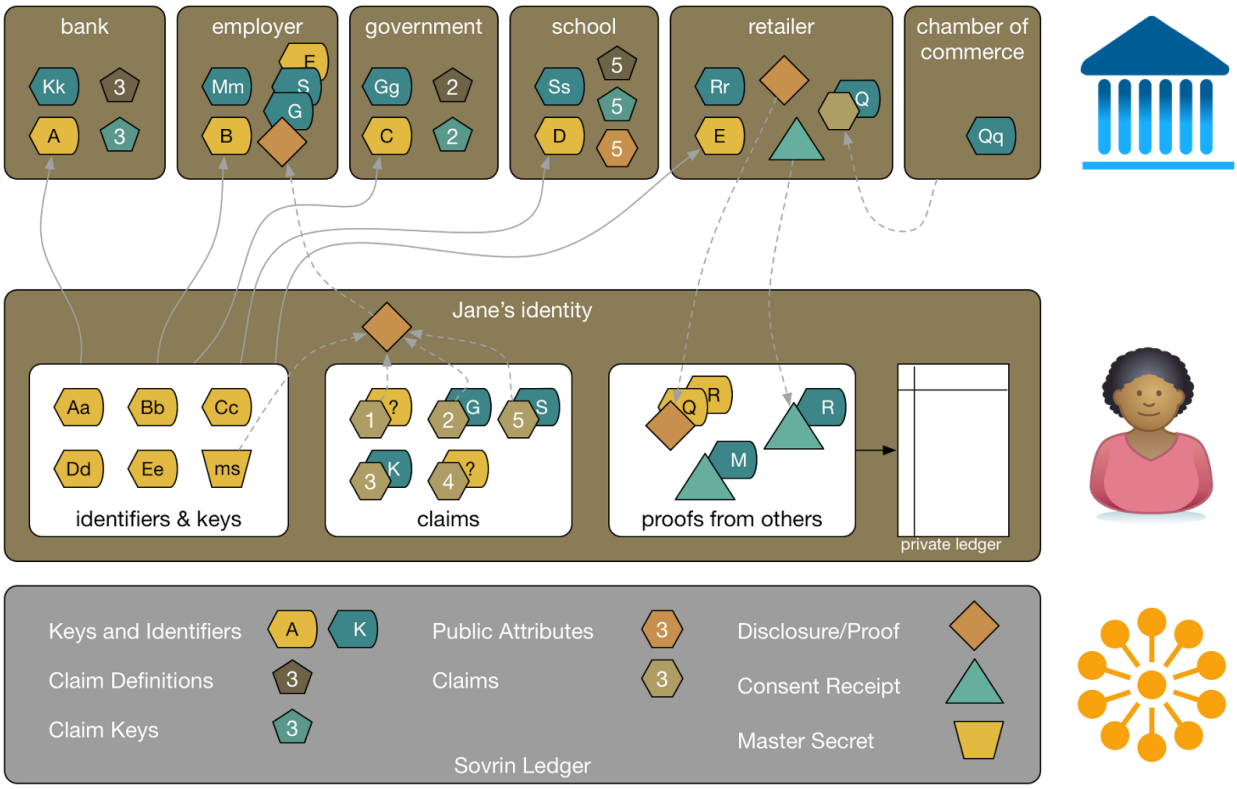

Jane's private ledger

One important purpose of the private ledger is to track the time and order that things are written. Hashes of the private ledger can be written to Sovrin's public ledger periodically to provide evidence of Jane's holdings. These are called anchors. Jane can thus generate audit proofs that any particular interaction happened at a particular time without disclosing the details of the interaction.

Jane may have a private ledger on her mobile phone, and in her browser. An agency could help her manage them and facilitate synchronization between her devices. Agencies are service providers that help Jane manage her identity. You can think of them as analogous to Internet Service Providers. They are substitutable for one another and Jane can choose a new agency without losing any of her identifiers, attributes, proofs, or claims. Because all of the information stored in Jane's private ledger is encrypted, there's a low security risk in using an agency as they never see Jane's transactions except in an encrypted state. Agencies are under legal contract to Sovrin Foundation to behave in specific ways to ensure Jane's security and privacy.

Advantages of Sovrin

We've see in the preceding examples a number of advantages that Sovrin holds over other identity systems.

- Public—everyone can use Sovrin. Both individuals and organizations will be relying parties and claim issuers.

- Permissioned—the Sovrin Foundation establishes rules for network nodes like stewards and agencies to hold them legally accountable for their actions.
- Reputation enhancing—Sovrin allows for the social proof of information people claim to be true, building reputation and thus enhancing trust.
- Trustworthy—identifiers and claims can be trusted because they are based on strong cryptography and governed by the Sovrin Foundation.
- Privacy enhancing—claims can be reused without risking correlation between identities. Sovrin reduces the hesitancy people might feel by enhancing privacy and thus reducing risk.
- Friction reducing—services can be integrated on the ledger, enhancing interactions between businesses

If you're interested in a deeper understanding of Sovrin and how it works, you should start by reading the Sovrin whitepaper The Inevitable Rise of Self-Sovereign Identity (PDF). There are also many other documents, including The Technical Foundation of Sovrin(PDF) and a Getting Started tutorial (PDF) on the Sovrin website.

Animation

An animated GIF showing the stages above can be found here:
http://www.windley.com/archives/2016/10/Sovrin-Animation.gif

How Do I Contact The Sovrin Foundation for More Information?

Please use the Contact Us page on sovrin.org